

Large deviations for the 3D dimer model

Nishant Chandgotia¹, Scott Sheffield², and Catherine Wolfram³

¹Tata Institute of Fundamental Research - Centre for Applicable Mathematics

²Institute for Advanced Study and Massachusetts Institute of Technology

³Massachusetts Institute of Technology

June 19, 2023

Abstract

In 2000, Cohn, Kenyon and Propp studied uniformly random perfect matchings of large induced subgraphs of \mathbb{Z}^2 (a.k.a. dimer configurations or domino tilings) and developed a large deviation theory for the associated *height functions*. We establish similar results for large induced subgraphs of \mathbb{Z}^3 . To formulate these results, recall that a perfect matching on a bipartite graph induces a *flow* that sends one unit of current from each even vertex to its odd partner. One can then subtract a “reference flow” to obtain a *divergence-free flow*. (On a planar graph, the curl-free dual of this flow is the height function gradient.)

We show that the flow induced by a *uniformly random* dimer configuration converges in law (when boundary conditions on a bounded $R \subset \mathbb{R}^3$ are controlled and the mesh size tends to zero) to the deterministic divergence-free flow g on R that maximizes

$$\int_R \text{ent}(g(x)) \, dx$$

given the boundary data, where $\text{ent}(s)$ is the maximal specific entropy obtained by an ergodic Gibbs measure with mean current s . The function ent is not known explicitly, but we prove that it is continuous and *strictly concave* on the octahedron \mathcal{O} of possible mean currents (except on the edges of \mathcal{O}) which implies (under reasonable boundary conditions) that the maximizer is uniquely determined. We further establish two versions of a large deviation principle, using the integral above to quantify how exponentially unlikely the discrete random flows are to approximate *other* deterministic flows.

The planar dimer model is mathematically rich and well-studied, but many of the most powerful tools do not seem readily adaptable to higher dimensions (e.g. Kasteleyn determinants, McShane-Whitney extensions, FKG inequalities, monotone couplings, Temperleyan bijections, perfect sampling algorithms, plaquette-flip connectivity, etc.) Our analysis begins with a smaller set of tools, which include Hall’s matching theorem, the ergodic theorem, non-intersecting-lattice-path formulations, and double-dimer cycle swaps. Several steps that are straightforward in 2D (such as the “patching together” of matchings on different regions) require interesting new techniques in 3D.

MSC2020 Subject Classifications: Primary 60F10, 82B20; Secondary 82B30. Keywords: Dimer tilings, domino tilings, perfect matchings, large deviations principle, variational principle, Gibbs measures, Hall’s matching theorem, chain swapping, local move connectedness.

Contents

1	Introduction	3
1.1	Overview	3
1.2	Two-dimensional background	4
1.3	Three-dimensional setup and simulations	6
1.4	Main results and methods	8
1.5	Three-dimensional history and pathology	16
1.6	Outline of paper	18
2	Preliminaries	20
2.1	Tilings and discrete vector fields	21
2.2	Measures on tilings and mean currents	22
2.3	Entropy	24
3	Local moves	25
3.1	Local moves in two dimensions	26
3.2	Local moves in three dimensions	26
3.3	Loop shift Markov chain for uniform sampling	30
3.4	Local move connectedness on the torus and k-Gibbs measures	31
4	Measures with boundary mean current	33
4.1	Review: EGMs with boundary mean current in two dimensions	34
4.2	EGMs with boundary mean current in three dimensions	35
5	Free-boundary tilings, asymptotic flows, and Wasserstein distance	42
5.1	Background on (generalized) Wasserstein distance	43
5.2	Wasserstein distance for flows	45
5.3	Main theorems	49
5.4	Boundary values of asymptotic flows	52
5.5	Boundary values of tiling flows	57
6	Patching	63
6.1	Hall's matching theorem and non-tileability	64
6.2	Discrete surfaces and minimal counterexamples	66
6.3	Statement of patching theorem and outline of proof	68
6.4	Discrete isoperimetric inequalities	74

6.5	Area growth of minimal monochromatic discrete surfaces	75
6.6	Tilings sampled from ergodic measures	77
6.7	Proof of the patching theorem	81
6.8	Corollaries for ergodic Gibbs measures	83
7	Properties of entropy	84
7.1	Entropy maximizers of a given mean current are Gibbs measures	85
7.2	Basic properties of ent	88
7.3	Flows for the double dimer model	89
7.4	Chain swapping	93
7.5	Strict concavity of ent and existence of EGMs of every mean current	104
7.6	Properties of Ent	106
8	Large deviation principles	111
8.1	Statement and set up: soft boundary LDP	113
8.2	Statement and set up: hard boundary LDP	116
8.3	Piecewise constant approximation	118
8.4	Existence of tiling approximations	121
8.5	Soft boundary lower bound	128
8.6	Generalized patching and hard boundary lower bound	130
8.7	Upper bounds	141
9	Open problems	145
	References	149
	Index of symbols	154
	Index of terms	156

1 Introduction

1.1 Overview

Let $G = (V, E)$ be a bipartite graph. A *dimer cover* (a.k.a. *perfect matching*) of G is a collection of edges so that every vertex is contained in exactly one edge. Throughout this paper, we will assume that G is an induced subgraph of \mathbb{Z}^d . We partition \mathbb{Z}^d into *even* vertices (the sum of whose coordinates is even) and *odd* vertices (the sum of whose coordinates is odd). By convention, we represent an (*a priori* undirected) edge e by (a, b) where a is even and b is odd.

We can also take a dual perspective, where each vertex $a = (a_1, a_2, \dots, a_d)$ is represented by the hypercube $[a_1 - \frac{1}{2}, a_1 + \frac{1}{2}] \times \dots \times [a_d - \frac{1}{2}, a_d + \frac{1}{2}]$ and each matched edge is represented by a “domino” which is the union of two adjacent hypercubes. The dominoes are 2×1 (or 1×2) boxes in 2D and $2 \times 1 \times 1$ (or $1 \times 2 \times 1$ or $1 \times 1 \times 2$) boxes in 3D. From this perspective, the perfect matchings corresponding to the subgraph induced by $V \subset \mathbb{Z}^d$ correspond to *domino tilings* of the region formed by the union of the corresponding cubes. The figure below illustrates a domino tiling of a two-dimensional region called the *Aztec diamond*. On the left, the domino corresponding to (a, b) is colored one of four colors, according to the value of the unit-length vector $b - a$. On the right, squares are colored by parity.




Figure 1: Tiling of an Aztec diamond and bipartite coloring of squares in \mathbb{Z}^2 .

In other words, every domino in the tiling on the left contains one square that is black (in the chessboard coloring on the right) and one that is white—and the color of a domino depends on whether its white square lies north, south, east or west of its black square.

Tilings with all dominoes oriented the same way are called *brickwork tilings*. There are four brickwork orientations in dimension 2—and $2d$ brickwork orientations in dimension d .




Figure 2: The four brickwork patterns in two dimensions.

A perfect matching τ of an induced subgraph R of \mathbb{Z}^d induces an *lattice flow* v_τ that sends one unit of current from every even vertex to the odd vertex it is matched to. If we subtract a “reference flow” (which sends a current of magnitude $1/2d$ from each even vertex to each of its $2d$ odd neighbors) we obtain a *divergence-free* flow f_τ . The main problem in this paper is to understand the behavior of the *random* divergence-flow f_τ that corresponds to a τ chosen uniformly from the set of tilings of a large region, subject to certain boundary conditions.

1.2 Two-dimensional background

In two dimensions, the divergence-free flow on \mathbb{Z}^2 described above has a dual flow on the dual graph (obtained by rotating each edge 90 degrees counterclockwise about its center)

that is a curl-free flow, and hence can be realized as the discrete gradient of some real-valued function defined on the vertices of the dual graph; see Section 2.1. This function (defined up to additive constant) is called the *height function* of the flow. Questions about the random flows associated to random perfect matchings can be equivalently formulated as questions about random height functions. For example, one can ask: when a tiling of a large region is chosen uniformly at random, what does the “typical” height function look like?

In 2000, Cohn, Kenyon and Propp studied domino tilings of a domain $R \subset \mathbb{R}^2$ like the one below, asking what happens in the limit as the mesh size tends to zero and the (appropriately rescaled) height function on the boundary converges to a limiting function [CKP01]. Note that given any tiling τ that covers R (e.g. τ could be one of the brickwork tilings) one can form a tileable region R_n by restricting to the tiles strictly contained in R , and the choice of R_n determines how the height function changes along the boundary.




Figure 3: An example of a fixed boundary region $R_n \subset R$ for the LDP in two dimensions.

Cohn, Kenyon and Propp showed that as the mesh size tends to zero (and the rescaled boundary heights converge to some function on ∂R) the random height function converges in probability to the unique continuum function u that (given the boundary values) minimizes the integral

$$\int \sigma(\nabla u(z)) dz, \tag{1}$$

where $\sigma(s)$ is the *surface tension* function, which means that $-\sigma(s)$ is the *specific entropy* (a.k.a. *entropy per vertex*) of any ergodic Gibbs measure of slope s . More generally, they established a theory of *large deviations* by showing how exponentially unlikely the random height function would be to concentrate near any other u . Earlier work studied this problem specifically for the *Aztec diamond*, see [CEP96] and [JPS98].

The proof in [CKP01] used ingredients from the scalar height function theory (McShane-Whitney extensions, monotone couplings, stochastic domination, etc.) and the Kasteleyn determinant representation (an exact formula for the entropy function) that do not appear readily adaptable to three dimensions.

The literature on the two-dimensional dimer model is quite large and we will not attempt a detailed survey here. Introductory overviews with additional references include e.g. [Ken09] and [Gor21].

We remark that fixing the asymptotic height function boundary values on ∂R is equivalent to fixing the asymptotic rate at which current flows through ∂R in the corresponding divergence-free flow. The latter interpretation is the one that extends most naturally to higher dimensions.

1.3 Three-dimensional setup and simulations

The goal of this paper is to extend [CKP01] to higher dimensions, where different tools are required. For simplicity and clarity, we focus on 3D, but we expect similar arguments to work in dimensions higher than three. (We discuss possible generalizations and open problems in Section 9.) Before we present our main results, we provide a few illustrations. The figure below illustrates a uniformly random tiling of a $10 \times 10 \times 10$ cube, with six colors corresponding to the six orientations. Next to it is the underlying black-and-white checkerboard grid. This figure and the others below were generated by a Monte-Carlo simulation (see Section 3.3) that is known to be mixing, but whose mixing rate is not known. It was run long enough that the pictures *appeared* to stabilize but we cannot quantify how close our samples are to being truly uniform. The efficient *exact sampling* algorithms that work in 2D do not have known analogs in 3D.




Figure 4: A dimer tiling of the $10 \times 10 \times 10$ cube and the bipartite coloring of the cubes in \mathbb{Z}^3 .

The figure below represents a random tiling τ of a region R called the *Aztec pyramid* (formed by stacking Aztec diamonds of width $2, 4, 6, \dots, 36$). Next to it is again the underlying black-and-white checkerboard coloring. Recall that (due to the reference flow) the divergence-free flow f_τ sends a $1/6$ unit of current through each square on the boundary ∂R . Such a square divides a cube inside R from a cube outside R . The flow is directed *into* R if the cube inside R is even, and *out of* R if the cube inside R is odd. Of the four triangular faces of the pyramid, two consist entirely of even cubes and the other two consist entirely of odd cubes. This means that f_τ current enters two opposite triangular faces at its maximal rate, and exits other two triangular faces at its maximal rate, while the net current through the bottom square face is zero (since on the lower boundary, the number of faces bounding even cubes in R equals the number of faces bounding odd cubes in R).




Figure 5: A dimer tiling of an Aztec pyramid and the bipartite coloring of the cubes in \mathbb{Z}^3 .

Below is a larger Aztec pyramid seen from above and from underneath. One can construct a computer animation showing the horizontal cross-sections one at a time. For the three large simulations shown here in Figures 6, 7 and 8, animations of the slices are available at <https://github.com/catwolfram/3d-dimers>. In these animations, it appears that each cross section has four “frozen” brickwork regions and a roughly circular “unfrozen” region, similar to the 2D Aztec diamond.




Figure 6: A tiling of a larger pyramid of Aztec diamonds, from the side and from below.

We now describe two larger labeled figures. Figure 7 illustrates a uniformly random tiling of the *Aztec octahedron* formed by gluing two Aztec pyramids along their square face. Four of the eight triangular faces of the octahedron contain only even cubes on their boundary, and the other four contain only odd cubes (and these alternate; distinct faces sharing an edge have opposite parity). In light of this, we can say that the current enters four of the faces at the maximum possible rate and exits the other four faces at the maximum possible rate. In simulations there appear to be twelve frozen regions (one for each *edge* of the octahedron) in which one of the six brickwork patterns dominates. (By contrast, tilings of the two dimensional Aztec diamond have four frozen regions, one for each *vertex* of the diamond.) Within each brickwork region, current flows at its maximum possible rate from one face (where it enters the octahedron) to an adjacent face (where it exits). Away from these brickwork regions, one sees a mix of colors, with a gradually varying density for each color. These are regions where the magnitude of the current flow is smaller, and the mean current appears to vary continuously across space.




Figure 7: Tiling of an Aztec octahedron.

Figure 8 illustrates a tiling of the *Aztec prism* (formed by stacking Aztec diamonds whose widths alternate between $2n$ and $2n + 2$). Again, each slice seems to be frozen outside of a roughly circular region. The width-alternation ensures that each of the four rectangular side faces of the prism has either only even faces or only odd faces exposed. Thus, current enters two of the opposite side faces at its maximal rate, and exits the other two at its maximal rate. The net current flowing through the top and bottom faces is zero. In this figure, and in all of the examples above, the distribution of domino colors in a subset of the tiled region determines the “mean direction of current flow” in that subset. We are interested in understanding what the “typical flow” looks like in the fine mesh limit.

1.4 Main results and methods

The main results of this paper will be two versions of a *large deviation principle* (LDP) for fine-mesh limits of uniformly random dimer tilings of compact regions $R \subset \mathbb{R}^3$, with some limiting boundary condition. The versions of the LDP we prove differ in how we treat the boundary conditions.




Figure 8: Tiling of an Aztec prism.

A large deviation principle is a result about a sequence of probability measures $(\rho_n)_{n \geq 1}$ which quantifies the probability of rare events at an exponential scale as $n \rightarrow \infty$. More precisely, a sequence of probability measures $(\rho_n)_{n \geq 1}$ on a topological space (X, \mathcal{B}) is said to satisfy a large deviation principle (LDP) with rate function I and speed v_n if $I : X \rightarrow [0, \infty)$ is a lower semicontinuous function, and for all Borel measurable sets B ,

$$-\inf_{x \in B^\circ} I(x) \leq \liminf_{n \rightarrow \infty} v_n^{-1} \log \rho_n(B) \leq \limsup_{n \rightarrow \infty} v_n^{-1} \log \rho_n(B) \leq -\inf_{x \in \bar{B}} I(x),$$

where B°, \bar{B} denote the interior and closure of B respectively. The rate function $I(x)$ is *good* if its sublevel sets $\{x : I(x) \leq a\}$ are compact. Implicit in this definition is a choice of topology \mathcal{B} on X . A large deviation principle for $(\rho_n)_{n \geq 1}$ implies that random samples from ρ_n are exponentially more likely to be near the minimizers of $I(\cdot)$ as $n \rightarrow \infty$. When I is good and has a unique minimizer, this means that random samples from ρ_n *concentrate* as $n \rightarrow \infty$ in the sense that if U is any neighborhood of the unique minimizer, then as $n \rightarrow \infty$ the probability $\rho_n(X \setminus U)$ tends to zero exponentially quickly in v_n . Good references for this subject include [DZ09] and [Var16].

To formulate the large deviation principle for dimer tilings, we associate to each tiling τ of \mathbb{Z}^3 a divergence-free discrete vector field. As mentioned above, we can define a flow v_τ that moves one unit of flow from the even endpoint to the odd endpoint of each e in τ . Subtracting a reference flow which sends $\frac{1}{6}$ flow along every edge from even to odd produces a divergence-free

discrete vector field:

$$f_\tau(e) = \begin{cases} +5/6 & e \in \tau \\ -1/6 & e \notin \tau. \end{cases}$$

We call this a *tiling flow*, and it will play an analogous role to the height function in two dimensions. The height function in two dimensions is a scalar potential function whose gradient is the *dual* of the tiling flow (i.e., the flow on the dual lattice obtained by rotating each edge 90 degrees counterclockwise about its center). See Section 2.1.

Tiling flows can also be *scaled* so that they are supported on $\frac{1}{n}\mathbb{Z}^3$ instead of \mathbb{Z}^3 . We scale them so that the total flow per edge is proportional to $\frac{1}{n^3}$, so that in the fine-mesh limit with $n \rightarrow \infty$, the total flow in a compact region in \mathbb{R}^3 converges to a finite value. The scaled tiling flows takes values $-\frac{1}{6n^3}$ and $\frac{5}{6n^3}$. A *scale n dimer tiling* is a dimer tiling of $\frac{1}{n}\mathbb{Z}^3$ (or a subset of it).

Fix a compact region $R \subset \mathbb{R}^3$ (which is sufficiently nice, e.g. the closure of a connected domain with piecewise smooth boundary). We define the *free-boundary tilings of R at scale n* to be tilings τ at scale n such that every point in R is contained in a tile in τ , and every tile has some intersection with R .




Figure 9: An example of a free-boundary tiling of R

We denote the corresponding *free-boundary tiling flows of R at scale n* by $TF_n(R)$. Note that $TF_n(R)$ is a finite set for all n . There is a signed *flux measure* on ∂R (supported on the points where edges of $\frac{1}{n}\mathbb{Z}^3$ cross ∂R) that encodes the net amount of flow directed into R . Since f_τ is divergence-free, the *total flux* through ∂R is zero. (See Definition 5.5.)

If τ is a *random* tiling of \mathbb{Z}^3 whose law is invariant under translations by even vectors, then one can define the *mean current per even vertex* to be $s = \mathbb{E}[\eta]$ where η is the vertex of \mathbb{Z}^3 matched to the origin by τ . Note that s lies in the set

$$\mathcal{O} = \{(s_1, s_2, s_3) \in \mathbb{R}^3 : |s_1| + |s_2| + |s_3| \leq 1\},$$

which is the convex hull of $\{(\pm 1, 0, 0), (0, \pm 1, 0), (0, 0, \pm 1)\}$, and which we call the *mean-current octahedron*. It indicates the expected total amount of current in v_τ (or equivalently f_τ) per even vertex; see Section 2.2. The vertices of \mathcal{O} arise for a random τ that is a.s. equal to one of the six brickwork tilings in three dimensions.

We define $AF(R)$ to be the space of measurable, divergence-free vector fields supported in R taking values in \mathcal{O} . The notation AF stands for *asymptotic flow* and is chosen because of the fact (formalized in Theorem 5.3.1) that these are precisely the flows that can arise as $n \rightarrow \infty$ limits of tiling flows on R , under a suitable topology.

The topology we use is the *weak topology* on the space of flows obtained by interpreting each vector component of the flow as a signed measure, see Section 5. This topology can also be generated by the *Wasserstein distance* for flows. Loosely speaking, two flows are considered Wasserstein close if one can be transformed into the other by moving, adding, and deleting flow with low “cost.” The cost is the amount of flow moved times the distance moved, plus the amount of flow added or deleted. The large deviation principles we prove use the same topology (weak topology, which is generated by Wasserstein distance) for the fine-mesh limits of random free-boundary tiling flows of R .

Remark 1.4.1. The 2D large deviation analysis in [CKP01] is based on a topology that at first glance looks different from ours, namely the topology in which two flows are considered close if their corresponding height functions h are close in L^∞ . However, it is not too hard to see that 1-Lipschitz functions on R (with fixed boundary values on ∂R) are L^∞ close if and only if their gradients (or equivalently the dual flows of their gradients) are Wasserstein close. We will not use or prove this fact here.

In three dimensions, we will also derive an LDP for random perfect matchings on finite regions approximating a continuum domain R , with boundary conditions chosen so that the flux through the boundary approximates a continuum boundary flow b , in a sense we will explain below. As in two dimensions, the rate function $I_b(\cdot)$ is a function of an asymptotic flow $g \in AF(R)$ and is (up to an additive constant) equal to

$$-\text{Ent}(g) = -\frac{1}{\text{Vol}(R)} \int_R \text{ent}(g(x)) \, dx, \quad (2)$$

where the additive constant is $C = \max_{f \in AF(R,b)} \text{Ent}(f)$. We interpret (2) as the three-dimensional analog of (1). The function $\text{ent}(s)$ is the maximal *specific entropy* of a measure with mean current $s \in \mathcal{O}$. It turns out that for every $s \in \mathcal{O}$ there exists an *ergodic Gibbs measure* of mean current s that achieves this maximal entropy $\text{ent}(s)$. This essentially follows from the strict concavity of ent when s is in the interior of \mathcal{O} , and can also be checked for $s \in \partial\mathcal{O}$. See Theorem 7.5.2 and Theorem 4.2.7.

To state our theorems, we need a way of fixing for each n a region R_n that approximates a continuum region R , such that boundary flow corresponding to R_n approximates a continuum boundary flow b on ∂R . The precise analog of the statement given in [CKP01] is not exactly true in 3D, due to subtleties related to the fact that in 3D the function ent can be nonzero even on the boundary of \mathcal{O} (see Section 4). To briefly illustrate what can go wrong, consider a finite region S tiled with only three types of tiles: north, east and up. Then every vertex $x = (x_1, x_2, x_3)$ with $x_1 + x_2 + x_3 = c$ (with c even) must be connected to a vertex y with $y_1 + y_2 + y_3 = c + 1$, and vice versa. The vertices with coordinate sums in $\{c, c + 1\}$ thus form a “slab” of points that are only connected to each other, and one can use Hall’s matching theorem to show that this must be the case for *any* tiling of S . Thus we can view S as a stack of two-dimensional slabs, where the tilings within the different slabs are independent of each other. These slabs could alternate between brickwork patterns (one slab has only east-going tiles, the next has only north-going tiles, the next has only up-going tiles, etc.) but they could also all be nonzero-entropy mixtures of the three tile types. Rescalings of both types of S might approximate the same continuum (R, b) , but the corresponding uniform-random-tiling measures would have very different entropy and very different large deviation behavior (see Example 8.2.6 and Section 4).

We will present two ways of formulating a theorem that *is* true in 3D. In the first approach we replace the *hard constraint* on the boundary behavior (where an R_n to be tiled is explicitly

specified for each n) with a *soft constraint* (where all scale n tilings that cover R are allowed, provided they induce boundary flows that are “sufficiently close” to the desired limiting flow b). This “soft constraint” LDP will apply to a fairly general set of pairs (R, b) . In the second approach we keep the hard constraints (i.e., the fixed R_n regions) but require (R, b) to be “flexible” in a sense that prevents the degenerate situation sketched above (where on the discrete level there might be slab boundaries that cannot be crossed by the tiles of *any* tiling of R_n). Precisely, we say (R, b) is *flexible* if for every $v \in R$ there exists an asymptotic flow g for (R, b) such that for some neighborhood $U \ni v$ we have $\overline{g(U)} \subset \text{Int}(\mathcal{O})$. Informally, this means there is no interior point near which g is *forced* to lie on $\partial\mathcal{O}$. For later purposes, we say that (R, b) is *semi-flexible* if for every $v \in R$ there exists an asymptotic flow g for (R, b) such that for some neighborhood $U \ni v$ the set $\overline{g(U)}$ is disjoint from the edges of \mathcal{O} . In other words, there is no interior point near which g is forced to lie on an edge of \mathcal{O} .

Using compactness of the space of flows, it is not hard to show that there exists g that minimizes (2). However it takes a bit more work to see whether such a g is unique. If ent were strictly concave everywhere, then $I_b(g_1)$ and $I_b(g_2)$ could not be minimal for distinct g_1 and g_2 (since the strict concavity would imply that $I_b(\frac{g_1+g_2}{2})$ was even smaller). The trouble is that ent is not strictly concave on the edges of \mathcal{O} (where it is constant) even though we will show that it is strictly concave everywhere else. In principle, there could still exist distinct minimizers g_1 and g_2 that (outside a set of measure zero) disagree *only* at points where they both take values on the same edge of \mathcal{O} . (See Problem 9.0.7.) We have not ruled out this possibility in general, but we can prove that (2) has a unique minimizer if (R, b) is semi-flexible. (See [Gor21, Proposition 7.10] for a 2D analog of this argument.) This in turn implies that the random flows in the corresponding LDP theorems *concentrate* near the unique minimizer g of I_b . If (R, b) is not semi-flexible then we call it *rigid*. We briefly summarize the conditions needed for each result in the following table before giving more precise statements.

(R, b)	SB LDP	I_b has unique minimizer	HB LDP
rigid	yes	not known	no
semi-flexible	yes	yes	no
flexible	yes	yes	yes

The results marked “no” in this table are provably not true in general. By taking limits of the “stack of slabs” regions discussed above, one can produce a semi-flexible (or rigid) pair (R, b) for which the hard boundary large deviation principle is false (for further discussion of this see Example 8.2.6).

Now, to introduce the soft boundary large deviation principle, we define the probability measure ρ_n to be the uniform measure on the space of flows in $TF_n(R)$ whose boundary values lie within Wasserstein distance θ_n of the desired limiting boundary flow, where $\theta_n \rightarrow 0$ as $n \rightarrow \infty$. We call θ_n the “threshold sequence” and it can be chosen arbitrarily provided that it does not tend to zero *too quickly* in a sense we explain later. A rough statement of our main theorem is the following.

Theorem 1.4.2 (See Theorem 8.1.6). *Let $R \subset \mathbb{R}^3$ be a compact region which is the closure of a connected domain, with piecewise smooth boundary ∂R . Let b be a boundary value for an asymptotic flow and let $(\theta_n)_{n \geq 1}$ be a (good enough) sequence of thresholds. Let ρ_n be uniform measure on $TF_n(R)$ conditioned on the boundary values being within θ_n of b .*

Then the measures $(\rho_n)_{n \geq 1}$ satisfy a large deviation principle in the Wasserstein topology on

flows with good rate function $I_b(\cdot)$ and speed $v_n = n^3 \text{Vol}(R)$, namely for any Wasserstein-measurable set A ,

$$-\inf_{g \in A^\circ} I_b(g) \leq \liminf_{n \rightarrow \infty} v_n^{-1} \log \rho_n(A) \leq \limsup_{n \rightarrow \infty} v_n^{-1} \log \rho_n(A) \leq -\inf_{g \in \bar{A}} I_b(g). \quad (3)$$

If g is an asymptotic flow, the rate function $I_b(g)$ is equal up to an additive constant to $-\text{Ent}(g)$. (Otherwise it is ∞ .)

Remark 1.4.3. The requirements for the region R are mild—for example, we do not require that R is simply connected. In this sense, the theorem can be viewed as extending both [CKP01] (simply connected 2D) and [Kuc22] (multiply connected 2D) to three dimensions. The requirement that ∂R be piecewise smooth is probably not necessary, but if the boundary of R is allowed to be too rough, the theorem statements one can make will depend more sensitively on how the boundary conditions are handled. For example, if the boundary of R has positive volume, then tilings that cover R may have volume-order more tiles than the tilings that approximate R “from within” and the extra tiles may contribute to the limiting entropy. If the boundary of R has infinite area, then the flux through the boundary may be an infinite signed measure, which would have to be defined more carefully. (For example, we could say that two flows that vanish outside of R have “equivalent boundary values” if their difference is a flow on all of \mathbb{R}^3 that is divergence-free in the distributional sense, and then let b denote an equivalence class.) For simplicity, we will focus on the piecewise smooth setting in this paper.

The distinction between soft and hard boundary conditions only substantially impacts one step of the proof: the argument that there exists a tiling (in the support of ρ_n) whose flow approximates a piecewise-constant flow that in turn approximates a given $g \in AF(R)$. Theorem 1.4.2 would still apply if the boundary conditions defining the ρ_n were specified in another way (ensuring convergence to (R, b) in the limit) as long as some version of this step could be implemented. We show using the *generalized patching theorem* (Theorem 8.6.2) that under the condition that (R, b) is flexible, this step can be implemented and a hard boundary LDP holds.

We say a boundary value is a *scale n tileable* if there exists a free-boundary tiling τ of R at scale n with that boundary value. If two tilings τ_1, τ_2 have the same boundary values on ∂R , then they are tilings of the same fixed region, so fixing a sequence of scale n tileable boundary values b_n is equivalent to fixing a sequence of *regions* R_n with boundary value b_n . A rough statement of the *hard boundary large deviation principle* is as follows.

Theorem 1.4.4 (Theorem 8.2.4). *Suppose that (R, b) is flexible and that $R_n \subset \frac{1}{n}\mathbb{Z}^3$ is a sequence of regions with tileable boundary values b_n on ∂R converging to b . Let $\bar{\rho}_n$ be uniform measure on dimer tilings of R_n . Then the measures $(\bar{\rho}_n)_{n \geq 1}$ satisfy a large deviation principle in the Wasserstein topology on flows with the same good rate function $I_b(\cdot)$ and speed $v_n = n^3 \text{Vol}(R)$ as the soft boundary measures ρ_n .*

It is straightforward to show that the (R, b) pairs obtained as fine-mesh limits of the “Aztec regions” above (Figures 6, 7, 8) are flexible, despite the fact that *typical* tilings appear (in simulations) to have frozen brickwork regions. (See Remark 8.2.3.)

Under the condition that (R, b) is *semi-flexible* (see further discussion in Definition 7.6.8) we show that Ent has a unique maximizer (Theorem 7.6.10). This together with some basic

topological results shows rigorously that *concentration* around a deterministic limit shape occurs, as we see in the simulations. This concentration holds for either the soft boundary measures ρ_n or the hard boundary measures $\bar{\rho}_n$.

Corollary 1.4.5 (See Corollary 8.1.9 and Corollary 8.2.7). *Assume that (R, b) is semi-flexible. For any $\epsilon > 0$, the probability that a uniformly random tiling flow on R at scale n (either sampled from ρ_n , i.e. with boundary flow conditioned to be in a shrinking interval around b , or sampled from $\bar{\rho}_n$ if (R, b) is flexible, i.e. with boundary flow conditioned to be a fixed value b_n converging to b) differs from the entropy maximizer with boundary value b by more than ϵ goes to 0 exponentially fast in n^3 as $n \rightarrow \infty$.*

The methods in this paper are substantially different from the methods used to prove the large deviation principle for dimer tilings in two dimensions. The two-dimensional dimer model is *integrable* or *exactly solvable* in the sense that one can derive a (beautiful) explicit formula for the specific entropy function analogous to our function ent , and this explicit formula is used in the large deviations proof. The three-dimensional dimer model is not known to be integrable in this way, so we rely on “softer” arguments. We comment on a few of these below.

One of the key ingredients which does have a 2D analog in [CKP01] is the *patching argument* (Theorem 6.3.5) which essentially states that if two tilings τ_1, τ_2 satisfy a requirement that they “asymptotically have the same mean current s ” for some $s \in \text{Int}(\mathcal{O})$, then we can cut out a bounded portion of τ_2 and patch it into an unbounded portion of τ_1 by tiling a thin annular region.




Figure 10: Schematic for the patching theorem.

For the hard boundary large deviation principle, we also prove a *generalized patching theorem* (Theorem 8.6.2), which says roughly that two tilings can be patched on a *general* annular region $R \setminus R'$ if they have the same boundary value b on ∂R and the inner one approximates a *flexible* flow $g \in AF(R, b)$.

Proving the patching theorems will be one of the more challenging aspects of this paper. The main input is Hall’s matching theorem (proved by Hall in 1935 [Hal35]) which gives us a criterion to check if a region (e.g. the annular region between the two tilings) is tileable by dimers. It turns out that the criterion we need to check can be phrased in terms of the existence of a discretized minimal surface and leads to an interesting sequence of constructions described for boxes in Section 6 and generalized in Section 8. These arguments are more involved than the 2D patching arguments in [CKP01], which rely on height functions and Lipschitz-extension theory. It is hard to summarize the argument without giving the details, but the following is a very rough attempt (which can be skimmed on a first read).

1. For each n , define an annular region A that we want to tile (which is roughly a scale n approximation of a fixed continuum annular region, with outer boundary conditions determined by τ_1 and inner by τ_2). Use Hall’s matching theorem to show that if A is non-tileable there must exist a “surface” dividing the cubes in A into two sub-regions such that 1) the cubes with faces on the surface are odd if they are in the first sub-region, even if in the second and 2) the first sub-region has more even than odd cubes overall.
2. Reduce to the case that the surface is in some sense a “minimal monochromatic surface” given its boundary, which touches both the inside and outside boundaries of A . (Here monochromatic means that all cubes on one side of the surface are odd.)
3. Use an argument involving isoperimetric bounds to show that such a surface must have at least a constant times n^2 faces when n is large.
4. Show that the even-odd imbalance in the first sub-region *cannot* be as large as it would need to be to provide a non-tileability proof. Do this by covering the first region by dominoes (from an tiling sampled from an ergodic measure in Section 6, or from a tiling that approximates a flow g whose existence is guaranteed by the flexible condition in Section 8) and use geometric considerations to show that there must be a lot of dominoes with only an odd cube in the first sub-region (including order n^2 in the middle of A) and relatively fewer dominoes with only an even cube in the first sub-region (using the ergodic theorem and the fact that both tilings approximate the same constant flow s , or in the generalized case by using Wasserstein distance bounds that apply near the boundary of A). Conclude that the first sub-region has at least as many odd cubes as even cubes, and hence does not prove non-tileability. This argument shows that there exists no surface that proves non-tileability and hence (by Hall’s matching theorem) the region is tileable.

Another key step in proving the main theorems is to derive properties of the entropy function ent despite not being able to compute it explicitly on all of \mathcal{O} . Instead, $\text{ent}(s)$ is defined abstractly as the maximum specific entropy $h(\cdot)$ of a measure with mean current s (see Section 2.3). From a straightforward adaptation of the classical variational principle of Lanford and Ruelle [LR69], it follows that $\text{ent}(s)$ is always realized by a *Gibbs measure* of mean current s (see Theorem 7.1.2, see also Section 2 for the definition of a Gibbs measure).

To prove strict concavity of ent on $\text{Int}(\mathcal{O})$ (Theorem 7.5.1), we note that a translation invariant measure μ with mean current s and with $h(\mu) = \text{ent}(s)$ must be a Gibbs measure, and we then use a variant of the cluster swapping technique used in [She05] to compare measures of different mean currents. We call this variant *chain swapping*. It is an operation on measures on *pairs* of dimer tilings (τ_1, τ_2) . From a pair of tilings (τ_1, τ_2) (sampled from μ), chain swapping constructs a pair of tilings (τ'_1, τ'_2) by “swapping” the tiles of τ_1 and τ_2 along some of the infinite paths in (τ_1, τ_2) with independent probability $1/2$ (or any probability $p \in (0, 1)$). We say that (τ'_1, τ'_2) is sampled from the *swapped measure* μ' . See Section 7.4 for a more detailed discussion of chain swapping, including Figure 29 for an example.

At a high level, chain swapping is an operation that allows us to take a coupling μ of measures μ_1, μ_2 on dimer tilings of mean currents s_1, s_2 , and construct a coupling μ' of two new measures μ'_1, μ'_2 on dimer tilings both of mean current $\frac{s_1+s_2}{2}$. We show that this operation preserves the total specific entropy (i.e. $h(\mu) = h(\mu')$) and ergodicity, but *breaks* the Gibbs property. More precisely, if μ_1, μ_2 are ergodic Gibbs measures of mean currents s_1, s_2 and $\frac{s_1+s_2}{2} \in \text{Int}(\mathcal{O})$,

then μ'_1, μ'_2 are *not* Gibbs, and hence do *not* have maximal entropy among measures of mean current $\frac{s_1+s_2}{2}$. The proof that the Gibbs property is broken under chain swapping requires very different techniques from those used in [She05].

Under the assumption that there exists an *ergodic* Gibbs measure μ_s of mean current s for any $s \in \mathcal{O}$, and that $\text{ent}(s) = h(\mu_s)$, strict concavity would follow easily: let $\mu_1 = \mu_{s_1}, \mu_2 = \mu_{s_2}$ and apply chain swapping to get new measures μ'_1, μ'_2 of mean current $\frac{s_1+s_2}{2}$. Since total entropy is preserved,

$$h(\mu'_1) + h(\mu'_2) = h(\mu_1) + h(\mu_2) = \text{ent}(s_1) + \text{ent}(s_2).$$

On the other hand, since μ'_1, μ'_2 are not Gibbs,

$$h(\mu'_1) + h(\mu'_2) < 2\text{ent}\left(\frac{s_1 + s_2}{2}\right),$$

which would complete the proof. A rigorous proof of the theorem is given in Section 7.5, and relies on casework based on ergodic decompositions as we do not know, a priori, that ergodic Gibbs measures of mean current s exist for all $s \in \mathcal{O}$. However it will then follow *from* strict concavity that this is true, and there exist ergodic Gibbs measures of all mean currents $s \in \mathcal{O}$ (Corollary 7.5.4).

The above is a discussion of ent on $\text{Int}(\mathcal{O})$, where no explicit formula is known. We remark that ent is explicitly computable when restricted to $\partial\mathcal{O}$, since it reduces to a two-dimensional problem (see Section 4).

1.5 Three-dimensional history and pathology

The three-dimensional model is fundamentally different from the two-dimensional version in many respects. To give one example, we recall that if τ and σ are distinct perfect matchings of \mathbb{Z}^2 that agree on all but finitely many edges, then one can construct a sequence $\tau = \tau_0, \tau_1, \tau_2, \dots, \tau_n = \sigma$ of perfect matchings such that for each k , the matchings τ_k and τ_{k-1} agree on all edges except those contained in a single 2×2 square — and on that square one of $\{\tau, \sigma\}$ has two parallel vertical edges and the other has two parallel horizontal edges [Thu90]. From the domino tiling point of view, we say that τ_{k-1} and τ_k differ by a *local move* that corresponds to rotating a pair of dominoes as shown below.



Figure 11: A local move or flip in 2D.

It turns out that the analogous statement is false in 3D. In fact, as we will explain in Section 3, there is *no* collection of local moves for which the analogous statement would be true in 3D. In 3D, one can construct (for any $K > 0$) a tiling τ of \mathbb{Z}^3 that is

1. *non-frozen* — i.e., there exists a tiling $\tau' \neq \tau$ that disagrees with τ on finitely many edges.
2. *locally frozen* to level K — i.e., there exists no tiling $\tau' \neq \tau$ that disagrees with τ on fewer than K edges.

To understand why this is the case, recall that in two dimensions, one can superimpose an arbitrary perfect matching with a brickwork matching to obtain a collection of non-intersecting left-to-right lattice paths as follows:




Figure 12: 2D non-intersecting paths.

There is an obvious bijection between non-intersecting path ensembles (as shown above) and dimer tilings (which is one way to deduce the integrability of the dimer model in two dimensions). Applying local moves corresponds to shifting these paths up and down locally. One can analogously superimpose a red three-dimensional matching with a black brickwork matching, to obtain an ensemble of left-to-right paths in three dimensions. But in this case the function that maps each “left endpoint” to the “right endpoint” on the same path may not be uniquely determined, as the following example shows. For clarity, the black and red edges that coincide with each other are not drawn—so both figures indicate a red perfect matching of \mathbb{Z}^3 that (when restricted to the box) consists only of right-going edges in the brickwork pattern (not shown) and a few non-right-going edges (shown together with the black right-going edges that share their endpoints).




Figure 13: Two examples of 3D non-intersecting paths with the same endpoints.

In general, there could be many different paths, and many ways to permute the wires from the left before plugging them in on the right. In the example above, the paths are “taut” in the sense that they have no freedom to “move locally” using local moves that change only, say, three or four edges at a time (and they can be extended to taut paths on all of \mathbb{Z}^3). In general, 3D path ensembles are not nicely ordered from top to bottom, and do not have the same lattice structure that 2D path ensembles enjoy. They can be braided in complicated ways.

Despite this complexity, various “local move connectedness” results for 3D tilings have been obtained. See, for example, the series of works by subsets of Freire, Klivans, Milet and Saldanha [MS14a, MS14b, Mil15, MS15, FKMS22, Sal22, Sal20, KS22, Sal21], the recent work [HLT23] by Hartarsky, Lichev, and Toninelli, and physics papers by Freedman, Hast-

ings, Nayak, Qi, and separately Bednik about topological invariants and so-called Hopfions [FHNQ11, Bed19a, Bed19b].

One of the basic observations is that even on box-shaped regions in 3D, one cannot transform any tiling to any other tiling with a sequence of *flips* (which swap two edges of a lattice square with the other two). There is a quantity associated to a tiling, called the *twist* (related to the linking number from knot theory) that is preserved by flip moves but changed by so-called *trit* moves, which involve removing three edges contained in the same $2 \times 2 \times 2$ cube and replacing them with three others, see below:




Figure 14: A *flip*, a *trit* and a flip-rigid configuration called a *hopfion*.

It remains open whether it is possible to connect any tiling of a rectangular box to any other using *both* the flip and trit moves shown above. It is still possible in 3D to generate random tilings of finite regions using Glauber dynamics (using an update algorithm that allows for the tiling to be modified along cycles of arbitrary length, see Section 3.3) but little is known about the rate of mixing (though bounds were given for another 3D tiling model in [RY00]).

Quitmann and Taggi have some additional important work on the 3D dimer model, which studies the behavior of loops formed by an independently sampled pair of dimer configurations [Tag22, QT22, QT23]. Among other things, they find that when one superimposes two independent random dimer tilings on an $n \times n \times n$ torus, the union of the tilings will typically contain cycles whose length has order n^3 .

Throughout this paper our basic physical intuition is that the 3D dimer model describes a steady current through a non-isotropic medium, and we are studying how the current varies in space. But we stress that papers like the one by Freedman et al [FHNQ11] have other field theoretic phenomena in mind (*topological excitations, Majorana fermions, Abelian anyons*, etc.) and we will not attempt to explain these interpretations here, though we will briefly mention a gauge theoretic interpretation of the dimer model in Problem 9.0.14.

Let us also remark that the literature on *related* topics is quite large, including (to give just a few examples) work on large deviations for graph homomorphisms $h : \mathbb{Z}^d \rightarrow \mathbb{Z}$ [KMT20], weakly non-planar dimer models [GMT20, GRT22], and a generalization of rhombus tilings to n dimensions [LMN01, Lam21, WMDB02].

1.6 Outline of paper

We establish notation and a few basic preliminaries in Section 2. We then illustrate the complexity of the 3D model with a brief discussion of the local move problem and related topics in Section 3.

In Section 4 we describe the ergodic Gibbs measures of *boundary mean current* (i.e., having mean current that lies on the boundary $\partial\mathcal{O}$, where \mathcal{O} is the octahedron of possible mean currents). Not all Gibbs measures with boundary mean current have zero entropy, but we can still compute the entropy function ent on $\partial\mathcal{O}$ by reducing it to a two-dimensional problem

(see Theorem 4.2.7).


Section 5 is a technical section where we present some of the function-theoretic preliminaries about flows. We define scaled tiling flows, the Wasserstein metric on flows for comparing them, and asymptotic flows (which we prove are the scaling limits of tiling flows in Theorem 5.3.1). We also define boundary values for both types of flows using a *trace operator* T , and show that T is uniformly continuous as a function of the flow (Theorem 5.5.7).

In Section 6, we deal with the fundamental problem of how one “patches together” regions of different tilings to form one perfect matching of a large region (Theorem 6.3.5). As noted above, the key tool is *Hall’s matching theorem*. We give an outline of the proof of Theorem 6.3.5 (in the “square annulus setting”) in Section 6.3 accompanied by a sequence of two dimensional pictures. In three dimensions, Hall’s matching theorem relates non-tileability to the existence of a certain type of minimal surface. The other key classical input in the proof of Theorem 6.3.5 is the isoperimetric inequality.

Section 7 concerns properties of the entropy functions ent and Ent , such as continuity, strict concavity, and uniqueness of maximizers. Since no exact formula for $\text{ent}(s)$ is known for mean currents s in the interior of \mathcal{O} this section involves interesting methods fairly different from dimension 2, in particular the chain swapping constructions in Section 7.4.

Section 8 finally ties together the ingredients of the previous sections to produce the two large deviation principles (Theorem 8.1.6 for soft boundary conditions and Theorem 8.2.4 for hard boundary conditions) which are our main results. Both of these are broken down into proving a lower bound on probabilities (Theorem 8.1.10 for soft boundary and Theorem 8.2.8 for hard boundary) and an upper bound (Theorem 8.1.11 for soft boundary and Theorem 8.2.9 for hard boundary). One of the slightly difficult parts of the paper is the explicit construction of a tiling flow approximating an asymptotic flow. This is a step in proving the lower bound which we call the “shining light” argument (Theorem 8.4.1). For the hard boundary lower bound, on top of this we also need a *generalized patching theorem* (Theorem 8.6.2) to show that any asymptotic flow can be approximated by a tiling of a fixed region. The proof of the generalized patching theorem is where we make use of the *flexible* condition on (R, b) in the hard boundary large deviation principle.

Several open problems are given in Section 9. See the chart below for a graphical representation of some of the dependencies and results that we highlight.



The results in orange boxes are stated using the Wasserstein metric for flows, and rely on many of its properties described in Section 5.

Acknowledgments. The authors have enjoyed and benefited from conversations with many dimer theory experts, including (but not limited to) Nathanaël Berestycki, Richard Kenyon, and Marianna Russkikh. The authors were partially supported by NSF grants DMS 1712862, DMS 2153742 and the first author was in addition supported by SERB Startup grant. The first author would like to thank Tom Meyerovitch for initiating his interest in the dimer model and introducing him to the flows associated with the model, and Kedar Damle and Piyush Srivastava for many helpful conversations regarding our simulations.

2 Preliminaries

As we mentioned earlier, it is sometimes convenient to represent a vertex of \mathbb{Z}^3 by the unit cube centered at that vertex, and to represent an edge $e = (a, b)$ of \mathbb{Z}^3 by the union of the two cubes centered at a and b (a domino). Both perspectives are useful for visualization, and we will use the terms *perfect matching* and *dimer tiling* somewhat interchangeably. We denote the space of dimer tilings of \mathbb{Z}^3 by Ω .

Recall that \mathbb{Z}^3 is a bipartite lattice, with bipartition into even points (where the coordinate sum is even) and odd points (where the coordinate sum is odd). In a dimer tiling of \mathbb{Z}^3 , there are six possible types of tiles corresponding to the six possible unit coordinate vectors. We denote the unit coordinate vectors by $\eta_1 = (1, 0, 0)$, $\eta_2 = (0, 1, 0)$, and $\eta_3 = (0, 0, 1)$. We denote the edge in \mathbb{Z}^3 connecting the origin to η_i by e_i and the edge connecting the origin to

$-\eta_i$ by $-e_i$.

2.1 Tilings and discrete vector fields

Given a perfect matching of \mathbb{Z}^3 , there is a natural way to associate a discrete, divergence-free vector field valued on oriented edges. We will call the flow corresponding to a tiling τ a *tiling flow*, denoted f_τ . Like height functions in two dimensions, tiling flows have well-defined scaling limits called *asymptotic flows* (which we describe in Section 5). Asymptotic flows capture the broad statistics of dimer tilings in a given region. Since our main results (e.g. our large deviation principle, analogous to [CKP01]) are related to the large scale statistics of dimer tilings, they are naturally formulated in terms of tiling flows.

Let E denote the set of edges in \mathbb{Z}^3 . A *discrete vector field* or *discrete flow* is a function from oriented edges of \mathbb{Z}^3 to the real numbers. Unless stated otherwise, we assume all edges are oriented from even to odd (flipping the orientation of the edge e reverses the sign of the discrete vector field on e). For a dimer tiling τ of \mathbb{Z}^3 , we associate a discrete vector field v_τ valued on the edges $e \in E$ defined by

$$v_\tau(e) = \begin{cases} +1 & \text{if } e \in \tau, \text{ oriented even to odd} \\ 0 & \text{if } e \notin \tau \end{cases} \quad (4)$$

We call v_τ the *pretiling flow*. Recall that by our definition of discrete vector fields, if e is oriented odd to even, we say that $v_\tau(e) = -1$. The divergence of a discrete vector field v is given by

$$\operatorname{div} v(x) = \sum_{e \ni x} v(e) \quad (5)$$

where the sum is over edges e oriented away from x (e.g. if x is even, the edges in the sum are oriented from even to odd, and the opposite if x is odd). From this definition, we compute that

$$\operatorname{div} v_\tau(x) = \begin{cases} +1 & \text{if } x \text{ is even} \\ -1 & \text{if } x \text{ is odd} \end{cases}$$

Therefore v_τ itself is not divergence-free, but the divergences don't depend on τ , so we can construct a divergence-free flow corresponding to a tiling τ by subtracting a fixed reference flow r . There are lots of reasonable choices for the reference flow. For simplicity and symmetry we choose:

$$r(e) = \frac{1}{6} \quad \text{for all edges } e \in E \text{ oriented from even to odd}$$

We can now define the *tiling flow* corresponding to a tiling τ of a region $R \subseteq \mathbb{Z}^3$.

Definition 2.1.1. Let τ be a dimer tiling of \mathbb{Z}^3 . The divergence-free, discrete vector field corresponding to τ is $f_\tau := v_\tau - r$. We call f_τ a *tiling flow*.

$$f_\tau(e) = \begin{cases} +5/6 & \text{if } e \in \tau \\ -1/6 & \text{if } e \notin \tau \end{cases}$$

If τ is a dimer tiling of a subgraph $G \subset \mathbb{Z}^3$, we define the tiling flow by restriction.

Remark 2.1.2. In dimension 2, the analogous definition of a tiling flow f_τ also works (in this case the reference flow is $1/4$ on all edges oriented from even to odd). For every discrete flow defined on edges (whose endpoints are vertices of \mathbb{Z}^2) there is a dual flow on dual edges (whose endpoints are faces of \mathbb{Z}^2) obtained by rotating each edge 90 degrees clockwise. If the original flow is divergence-free, then the dual flow is curl-free and is hence equal to the gradient of a function (this function is called the *height function* or *scalar potential*). It is also worth noting that there is an analog of the height function in three dimensions. Namely, since f_τ is a divergence-free flow it can be written as the curl of another flow, that is, $f_\tau = \nabla \times A$, where A is a so-called *vector potential* which is defined modulo the addition of a curl-free flow. However the set of vector potentials A is more complicated than the set of height functions (it does not have a similar lattice structure, the potentials are not uniquely defined, etc.) and is not as useful for our purposes as height functions are in two dimensions. Because of that, we do not work with the vector potential in this paper, and instead just work with the tiling flow f_τ itself.

A pair of dimer tilings $(\tau_1, \tau_2) \in \Omega \times \Omega$ is called a *double dimer tiling* or *double dimer configuration*. The double dimer model is a model of independent interest, but we mention it because it will be a tool for comparing dimer tilings. This will be used in Section 3, Section 8, and substantially in Section 7.

There is a natural way to associate a divergence-free discrete flow to a double dimer configuration, namely for e an edge oriented from even to odd,

$$f_{(\tau_1, \tau_2)}(e) = f_{\tau_1}(e) - f_{\tau_2}(e) = v_{\tau_1}(e) - v_{\tau_2}(e) = \begin{cases} 1 & \text{if } e \in \tau_1 \setminus \tau_2 \\ -1 & \text{if } e \in \tau_2 \setminus \tau_1 \\ 0 & \text{if } e \in \tau_1 \cap \tau_2 \text{ or if } e \notin \tau_1 \cup \tau_2. \end{cases} \quad (6)$$

Unlike the tiling flow for a single tiling, the flow associated with a double dimer configuration (τ_1, τ_2) does not determine (τ_1, τ_2) , since it does not specify the tiles in $\tau_1 \cap \tau_2$. However, the collection of loops formed by $\tau_1 \cup \tau_2$ (including the double tiles) and the flow $f_{(\tau_1, \tau_2)}$ together determine (τ_1, τ_2) . See Section 7.3 for more about double dimer flows.

2.2 Measures on tilings and mean currents

Recall that Ω denotes the set of dimer tilings of \mathbb{Z}^3 . The group \mathbb{Z}^3 acts naturally on Ω by translations, namely given $x \in \mathbb{Z}^3$ and $\tau \in \Omega$, $\tau + x$ is the tiling where $(a, b) \in \tau$ if and only if $(a + x, b + x) \in \tau + x$. There is natural topology on Ω induced by viewing it as a subset of $\{0, 1\}^E$ and giving the latter the product topology over the discrete set $\{0, 1\}$ (recall that E denotes the edges of \mathbb{Z}^3). This makes Ω a compact metrizable space and the translation action on it continuous.

Let $\mathbb{Z}_{\text{even}}^3$ denote the set of even vertices in \mathbb{Z}^3 . We define $\mathcal{P}(\Omega) = \mathcal{P}$ to be the space of Borel probability measures on Ω invariant under the action of $\mathbb{Z}_{\text{even}}^3$.

To explain why we look at $\mathbb{Z}_{\text{even}}^3$ -invariant measures instead of \mathbb{Z}^3 -invariant measures, recall that \mathbb{Z}^3 is a bipartite lattice, with bipartition consisting of even points and odd points. In the interpretation of a dimer tiling as a flow from in Section 2.1, the sign of the flow on an edge oriented parallel to $(1, 0, 0)$ (for example) depends on whether the edge starts at an even point or an odd point. E.g. consider the tiling

$$\tau = \{(x, x + (1, 0, 0)) : x \in \mathbb{Z}^3 \text{ is even}\}.$$

The flow associated to τ moves current (on average) in the direction $(1, 0, 0)$, while the flow for $\tau + (1, 0, 0)$ moves current (on average) in the direction $(-1, 0, 0)$. We want our measures to be invariant under an action that preserves the asymptotic direction of the flow associated to a tiling, and this is why we consider $\mathbb{Z}_{\text{even}}^3$ -invariant measures instead of \mathbb{Z}^3 -invariant measures.

The *ergodic measures* \mathcal{P}_e are the extreme points of the convex set \mathcal{P} . A good reference for basic ergodic theory suitable for our purposes is [Kel98]. Any invariant measure $\mu \in \mathcal{P}$ can be decomposed in terms of ergodic measures, i.e. there exists a measure w_μ on \mathcal{P}_e such that

$$\mu = \int_{\mathcal{P}_e} \nu \, dw_\mu(\nu).$$

The measures ν in the support of w_μ are called the *ergodic components* of μ . Sampling from μ can be viewed as first sampling an ergodic component ν from w_μ and then sampling from ν .

We will also frequently make use of the so-called *uniform Gibbs measures* on Ω defined as follows: a measure $\mu \in \mathcal{P}$ is a *uniform Gibbs measure* if for any finite set $R \subset \mathbb{Z}^3$, we can say that *given* that τ contains no edges that cross the boundary of R , and *given* the tiling τ induces on $\mathbb{Z}^3 \setminus R$, the μ *conditional law* of the restriction of τ to R is the uniform measure on dimer tilings of R . We will see in the next section that the measures that maximize specific entropy are uniform Gibbs measures. Throughout the rest of the paper, we refer to uniform Gibbs measures simply as *Gibbs measures*.

A useful reference for Gibbs measures is [Geo11]. We denote the set of $\mathbb{Z}_{\text{even}}^3$ -invariant Gibbs measures by \mathcal{P}_G and the set of ergodic Gibbs measures (EGMs) by $\mathcal{P}_{G,e}$. A useful fact throughout is that the ergodic components of $\mathbb{Z}_{\text{even}}^3$ -invariant Gibbs measures are themselves $\mathbb{Z}_{\text{even}}^3$ -invariant Gibbs measures.

Proposition 2.2.1. [Geo11, Theorem 14.15] *The ergodic components of an invariant Gibbs measure are ergodic Gibbs measures almost surely.*

We remark that the analogous constructions work for *weighted Gibbs measures*. For example, one may assign weights $a_1, a_2, a_3, a_4, a_5, a_6$ to the six possible tile orientations. A $\mathbb{Z}_{\text{even}}^3$ -invariant Gibbs measure μ with these weights is a measure where for any finite set R , the conditional law of μ given a tiling of $\mathbb{Z}^3 \setminus R$ is the one in which each tiling of R has probability proportional to $\prod_{i=1}^6 a_i^{N_i}$, where N_i is the number of tiles of weight a_i . We expect that our main results could be extended to weighted dimer models (and perhaps also models with weights that vary by location in a periodic way as in [She05, KOS06]) but for simplicity we focus on the unweighted case here.

A key invariant of a $\mathbb{Z}_{\text{even}}^3$ -invariant measure is a quantity called the *mean current* which (as mentioned in the introduction) represents the expected current flow *per even vertex*. This is a generalization of the notion of *height function slope* from two dimensions. Recall that η_1, η_2, η_3 denote the standard basis for \mathbb{Z}^3 , the edge connecting the origin with η_i is denoted by e_i , and the edge connecting the origin with $-\eta_i$ is denoted by $-e_i$.

Definition 2.2.2. The *mean current* of a measure $\mu \in \mathcal{P}$, denoted $s(\mu)$, is an element of \mathbb{R}^3 such that its i^{th} -coordinate is

$$(s(\mu))_i = \mu(e_i \in \tau) - \mu(-e_i \in \tau).$$

Note that the mean current is an affine and continuous function of the measure. The mean current is invariant under the action of $\mathbb{Z}_{\text{even}}^3$ and takes values in

$$\mathcal{O} = \{s \in \mathbb{R}^3 : |s_1| + |s_2| + |s_3| \leq 1\}$$

which we call *the mean current octahedron*.

There are a few other useful formulations of the mean current. We define the function $s_0 : \Omega \rightarrow \mathbb{R}^3$ to be the direction of the tile at the origin in τ . Then the mean current can be computed as an expected value of s_0 :

$$s(\mu) = \int_{\Omega} s_0(\tau) d\mu(\tau). \quad (7)$$

Similarly let $\Lambda_n = [-n, n]^3$, and let $\text{even}(\Lambda_n)$ denote the even points in Λ_n . We define the function

$$s_n(\tau) = \frac{1}{\text{even}(\Lambda_n)} \sum_{x \in \text{even}(\Lambda_n)} s_0(\tau + x). \quad (8)$$

The function $s_n(\tau)$ measures the average tile direction of τ in the box Λ_n . By $\mathbb{Z}_{\text{even}}^3$ -invariance,

$$s(\mu) = \int_{\Omega} s_n(\tau) d\mu(\tau). \quad (9)$$

We let \mathcal{P}^s denote the space of $\mathbb{Z}_{\text{even}}^3$ -invariant probability with mean current s . Adding the subscripts G and e will denote whether the measure is a Gibbs measure and whether it is ergodic with respect to the $\mathbb{Z}_{\text{even}}^3$ action.

2.3 Entropy

As is common in statistical physics models, entropy plays an important role in the large deviation principle for dimer tilings in 3D. There are a few different functions that we refer to as “entropy” (of a probability measure with finite or infinite support, of a mean current, of an asymptotic flow). Here we give some definitions and explain how these notions of entropy are related to each other. The primary reference for this section is also [Geo11].

For a probability measure ν with finite support S , its *Shannon entropy*, denoted $H(\nu)$, is

$$H(\nu) = - \sum_{\sigma \in S} \nu(\sigma) \log \nu(\sigma).$$

For a $\mathbb{Z}_{\text{even}}^3$ -invariant probability measure μ with infinite support, we can define the *specific entropy* of μ as a limit of Shannon entropy per site. Given a finite region $\Lambda \subset \mathbb{Z}^3$, let $\Omega(\Lambda)$ denote the dimer tilings of Λ (i.e. tilings of \mathbb{Z}^3 restricted to Λ , so tiles are allowed to have one cube outside Λ). For $\sigma \in \Omega(\Lambda)$, define

$$X(\sigma) = \{\tilde{\sigma} \in \Omega : \tilde{\sigma}|_{\Lambda} = \sigma\}$$

and then

$$H_{\Lambda}(\mu) := - \sum_{\sigma \in \Omega(\Lambda)} \mu(X(\sigma)) \log \mu(X(\sigma)).$$

Let $\Lambda_n = [-n, n]^3$ be a sequence of growing cubes. If μ is a $\mathbb{Z}_{\text{even}}^3$ -invariant probability measure on Ω , the *specific entropy* of μ , denoted $h(\mu)$, is

$$h(\mu) := \lim_{n \rightarrow \infty} |\Lambda_n|^{-1} H_{\Lambda_n}(\mu).$$

This limit exists because the terms form a subadditive sequence. In fact, one can also show that

$$h(\mu) = \inf_{\Lambda \in \mathcal{S}} |\Lambda|^{-1} H_{\Lambda}(\mu),$$

where \mathcal{S} is the set of all possible finite regions in \mathbb{Z}^3 . See [Geo11, Theorem 15.12]. As a function of the measure, $h(\cdot)$ is affine and upper semicontinuous [Geo11, Proposition 15.14].

The reason that Gibbs measures (introduced in the previous section) play a special role in our study is the *variational principle* which says that a measure $\mu \in \mathcal{P}$ maximizes $h(\cdot)$ if and only if μ is a Gibbs measure. This is a classical result going back to [LR69], see [Geo11, Theorem 15.39] for exposition.

The *local* or *mean-current entropy function* $\text{ent} : \mathcal{O} \rightarrow \mathbb{R}$ is defined

$$\text{ent}(s) = \max_{\mu \in \mathcal{P}^s} h(\mu).$$

This function is the main focus of Section 7, where we show it has a number of useful properties (continuity, concavity) and show that the maximum is always realized by an ergodic Gibbs measure of mean current s . In Theorem 4.2.7 we compute its restriction to $\partial\mathcal{O}$ by relating it to the analogous local entropy function for lozenge tilings in two dimensions.

We conclude this section with one more use of the term entropy. In Section 5, we will show that the “fine-mesh limits” of rescaled tiling flows are precisely the measurable vector fields we call *asymptotic flows*. Asymptotic flows are valued in \mathcal{O} and supported on some compact region R . The *entropy of an asymptotic flow* g can then be defined as

$$\text{Ent}(g) = \frac{1}{\text{Vol}(R)} \int_R \text{ent}(g(x)) \, dx.$$

3 Local moves

A number of the papers about the 3D dimer model are about local moves. Here we present some simple examples, briefly review the literature, and explain why local move connectedness fails for the torus in dimensions $d > 2$. Most of the ideas in this section are already known, but we include a few elementary observations we have not seen articulated elsewhere.

This section can be skipped on a first read, since the results are not essential for the rest of the paper. However, it is useful for understanding some of the ways that the $d = 2$ problem differs from the $d = 3$ problem (e.g., why the Kasteleyn determinant approach to computing entropy does not work in the same way) and also what makes $d = 3$ different from $d > 3$ (e.g., the integer-valued twist function is indexed by \mathbb{Z} when $d = 3$ and by $\mathbb{Z}/2\mathbb{Z}$ when $d > 3$). This section will also explain how the figures in the introduction were generated.




Figure 15: A local move or flip in two dimensions.




Figure 16: (1) an example of a sequence of local moves transforming one tiling into another and (2) the collection of cycles from overlaying the first and last tilings in this sequence.

3.1 Local moves in two dimensions

In two dimensions, a *local move* or *flip* is the operation of choosing a pair of parallel dimers in the tiling, and switching them out for the other pair. See Figures 15 and 16. Let R be a subgraph of \mathbb{Z}^2 and let $\mathcal{T}(R)$ be a graph on the set of dimer tilings of R where two tilings τ and τ' are connected by an edge if they differ by a single flip. It is shown using height functions in [Thu90] that if $R \subset \mathbb{Z}^2$ is simply connected and finite, then any two dimer tilings of R differ by a finite sequence of flips. In other words, $\mathcal{T}(R)$ is a connected graph.

Local move connectedness in the 2D dimer model means that it is possible to probe all tilings of a region using simple local updates, and this is useful for both theoretical and practical purposes. It means that uniformly random dimer tilings in 2 dimensions can be simulated using Markov chain Monte Carlo methods called *Glauber dynamics*. For the 2D dimer model, one can give an explicit polynomial bound on the mixing time of this algorithm [RT00].

3.2 Local moves in three dimensions

The same local moves (flips) make sense for the 3D dimer model, but local move connectedness with these manifestly fails, even for very small regions. There is a simple counterexample on the $3 \times 3 \times 2$ box which is called a *hopfion* in the physics literature (see Figure 17).

There is a series of papers by Fiere, Milet, Klivans, and Saldanha studying local move connectedness in dimension three under flips and *trits*, a new local move in three dimensions




Figure 17: A flip, a trit, and a flip-rigid configuration called a *hopfion*. The hopfion has no parallel pairs of tiles, so it is not connected under flips to any other tiling of the $3 \times 3 \times 2$ box.

involving three tiles (see Figure 17). In [MS15, MS14b] they show that any two tilings of a region of the form $D \times [0, 1]$ where D is simply connected and planar are connected under flips and trits. In subsequent works [MS14a, Sal22, Mil15, FKMS22, Sal21] they introduce and study an invariant called the *twist*, related to the linking number or writhing number. We will present below a brief and informal overview of the various ways the twist is defined and how it is related to a linking number. More detailed exposition is found e.g. in [Sal20] or the references above.

Given two distinct smooth curves $\gamma_1, \gamma_2 : S^1 \rightarrow \mathbb{R}^3$ embedded in \mathbb{R}^3 , one can compute their integer-valued *linking number* $L(\gamma_1, \gamma_2)$ by projecting them to a generic plane and summing the signatures of the crossings. (Recall that signature of a crossing of two oriented paths is 1 or -1 depending on whether the upper curve crosses the lower curve from right to left or left to right, when the bottom curve is viewed as being oriented from down to up.) It is a standard result that this number is independent of the plane one projects onto, see e.g. [Ada94, pages 20-21]. (The idea is to show that any one projection can be transformed into another by a sequence of Reidemeister moves, and that these moves preserve the linking number.) The linking number can also be computed with an integral formula: if r_1, r_2 are parametrizations of γ_1, γ_2 , then

$$L(\gamma_1, \gamma_2) = \frac{1}{4\pi} \oint_{\gamma_1} \oint_{\gamma_2} \frac{r_1 - r_2}{|r_1 - r_2|^3} dr_1 \times dr_2.$$

Informally, this represents the line integral along γ_1 of the magnetic field generated by a steady current through γ_2 . One can analogously compute a “linking number” of a pair of tilings in a box by summing crossings. Namely, imagine that each edge in the matching is extended $\epsilon > 0$ units in either direction. Then the crossing number is obtained by flattening these extended edges to a horizontal plane and summing the signatures of the crossings. To be more precise, we say two edges (a, b) and (c, d) constitute a crossing if their orientations are both orthogonal to the vertical (third-coordinate) direction and orthogonal to each other *and* one of the endpoints of (a, b) differs from one of the endpoints of (c, d) in the vertical coordinate and in no other coordinate. This is the same as an ordinary crossing if we assume each edge is extended ϵ units beyond its endpoints, and the sign of the crossing is defined in the usual way. We can define the linking of τ_1 and τ_2 to be the signed sum $L(\tau_1, \tau_2)$ of all crossings involving a tile in τ_1 and a tile in τ_2 . This is a quadratic form, and the twist of a tiling τ is defined by $T(\tau) = \frac{1}{4}L(\tau, \tau)$. This decomposes as a sum over pairs of horizontal tiles in vertical columns. For reasonable regions (i.e., $D \times [1, N]$, where $D \subset \mathbb{Z}^2$ is simply connected and N is even), the twist is integer-valued despite the $\frac{1}{4}$ and is independent of the direction for the orthogonal projection [MS14a, Proposition 6.4]. Within a rectangular box, one can easily show that trits increment the twist and flips leave the twist unchanged (in fact this holds for any region of the form $D \times [1, N]$, [MS14a, Theorem 1]). There are also examples of tilings with twist $T(\tau) = 0$ that are not connected under flips alone ([Sal21, Figure 7]), meaning that $T(\tau) = T(\sigma)$ does not imply that τ, σ are connected under flips.

Simple questions about local move connectedness under flips and trits still remain open, for example it is not known whether all tilings of an $M \times N \times L$ box are connected under flips and trits when $M, N, L > 2$ (see Problem 9.0.1). See [MS] for an enumeration of all tilings of the $4 \times 4 \times 4$ box, which shows explicitly that all tilings of this region are connected under flips and trits.

In dimensions $d > 3$, Klivans and Saldanha [KS22] show that the twist is valued in $\mathbb{Z}/2$. In dimension $d = 4$, even tilings of the $2 \times 2 \times 2 \times 2$ box fail to be connected under flips (see

[KS22, Example 2.2]). They also show that tilings within certain larger boxes are “almost” connected under flips, i.e. they can be connected if the boxes are extended in some way.

The works of Friere, Klivans, Milet and Saldanha rely mostly on geometric and algebraic constructions to study local move problems in dimensions $d \geq 3$, but the recent work [HLT23] by Hartarsky, Lichev, and Toninelli (written concurrently with this paper) makes progress using purely combinatorial arguments. In particular it follows from their results that any tiling of a rectangular box in \mathbb{Z}^3 which is tileable by dimers admits at least one flip or trit [HLT23, Theorem 3], providing a partial answer to Problem 9.0.1 in Section 9.

In fact, [HLT23, Theorem 3] is a statement that holds for any dimension $d \geq 2$. It states that any tiling τ of a rectangular box in \mathbb{Z}^d of dimensions (n_1, \dots, n_d) which is tileable by dimers contains a copy of $[0, 1]^d$ such that τ restricted to this copy of $[0, 1]^d$ contains at least $2^{d-2} + 1$ dimers. Specialized to the case $d = 3$, this means that there is a copy of $[0, 1]^3$ which completely contains at least three dimers from τ , and the only way this can happen is if $[0, 1]^3$ contains tiles which make up a flip or a trit in τ . The main idea of the proof is a clever but simple counting argument. Following the ideas in [HLT23], we present a slight modification of their proof specialized to the $d = 3$ case, with the aim of just showing the flip/trit result.

Proposition 3.2.1 ([HLT23]). *Let $R = [1, n_1] \times [1, n_2] \times [1, n_3] \subset \mathbb{Z}^3$ with $n_1, n_2, n_3 \geq 2$ and $n_1 n_2 n_3$ even. Any tiling τ of R admits at least one flip or trit.*

Proof. Fix a tiling τ of R . We view τ as a tiling of the torus with the same dimensions (i.e., τ is a tiling of the torus such that no dimers cross the identifications). On one hand, τ contains $n_1 n_2 n_3 / 2$ tiles, and each tile is contained in exactly four translates of $[0, 1]^3$. On the other hand, there are $n_1 n_2 n_3$ possible choices of translates of the unit cube in the torus, so the average number of tiles per unit cube is 2.

If a unit cube contains an above-average number of tiles from τ , it contains at least three tiles. If this unit cube is in the interior of R , or is cut in half by only one face of R , then since the tiles in τ do not cross the identifications, this implies there is a flip or trit in τ as a tiling of R . The result then follows by showing that the unit cubes which are cut into four pieces along the edges (or eight pieces at the corner) by the identifications of the torus contain a below-average number of tiles from τ .

The number of such “edge unit cubes” is $(n_1 - 1) + (n_2 - 1) + (n_3 - 1) + 1 = n_1 + n_2 + n_3 - 2$. Any dimer contained in an edge unit cube must be contained along one of the edges around R . The number of vertices in the edges around R is $4(n_1 + n_2 + n_3) - 16$ (there are 8 corners, but each one is contained in three edges), hence the maximum number of dimers contained in this region is $2n_1 + 2n_2 + 2n_3 - 8$. Given this, the average number of dimers in τ per edge unit cube is bounded by

$$\frac{2n_1 + 2n_2 + 2n_3 - 8}{n_1 + n_2 + n_3 - 2} < 2.$$

Therefore there must be a non-edge unit cube containing at least three tiles from τ , which completes the proof. \square

For the hypercube $[0, 1]^d \subset \mathbb{Z}^d$, Hartarsky, Lichev, and Toninelli also show that for $d \geq 3$, the connected components of the graph on dimer configurations of $[0, 1]^d$ under local moves of length up to $d - 1$ (here the trit is a move of length three and the flip is a move of length two) have size exponential in d [HLT23, Theorem 5], and that for $d \geq 2$, any two dimer tilings of

$[0, 1]^d$ are connected by a sequence of moves of length $\leq 2(d-1)$ [HLT23, Theorem 6]. For $[0, n]^d \subset \mathbb{Z}^d$, $d \geq 2$, n odd, they show that the diameter of the graph on dimer configurations of $[0, n]^d$ under local moves of length $\leq \ell$ is at least $n^{d-1}(n^2-1)/(6\ell^2)$ [HLT23, Theorem 7].

Flip connectedness has also independently been studied in the physics literature, from the perspective of looking for “topological invariants” preserved by flips. In [FHNQ11] the authors define a “Hopf number” for dimer tilings of \mathbb{Z}^d valued in $\pi_d(S^{d-1})$ which is invariant under flips. The hopfion (see Figure 17) has Hopf number ± 1 (depending on its orientation). This construction works for any dimension $d \geq 2$. The fact that $\pi_2(S^1) = 0$ corresponds to no obstruction to connectedness under flips, and $\pi_3(S^2) = \mathbb{Z}$ corresponds to there being at least countably many connected components under flips in dimension 3. For all $d > 3$, $\pi_d(S^{d-1}) = \mathbb{Z}/2$, implying at least two connected components under flips.

In [Bed19a] it is shown in examples that the Hopf number from [FHNQ11] can be computed using discrete versions of Cherns-Simon integral formulas for the Hopf number applied to a version of the tiling flow and its vector potential. See also [Bed19b].

Remark 3.2.2. The failure of local move connectedness in three dimensions is intimately related to the failure of (at least a straightforward generalization) of Kasteleyn theory.

In two dimensions, the partition function for dimer tilings of a simply connected planar graph can be computed as the Pfaffian of an adjacency matrix of the directed graph with appropriate weights (this can also be done with a determinant when the graph is bipartite). Recall that if $M = (m_{ij})$ is an $2n \times 2n$ skew-symmetric matrix,

$$\text{Pf}(M) = \frac{1}{2^n n!} \sum_{\sigma \in S_{2n}} \text{sign}(\sigma) \prod_{i=1}^n m_{\sigma(2i-1), \sigma(2i)}.$$

There are two key observations in two dimensions. First, the weights can be chosen so that the term is ± 1 if and only if the pairing $\{\sigma(2i-1), \sigma(2i)\}_{1 \leq i \leq n}$ corresponds to a dimer tilings and otherwise it is 0. By this, it is clear that the partition function can be computed as a *permanent* (i.e., like the above without the sign terms). The second key observation, which is why this reduces to a Pfaffian computation, is that the weights can be chosen so that applying a flip does not change the sign of the term. From here, flip connectedness in two dimensions shows that the Pfaffian is counting tilings.

In three dimensions it is still possible to choose weights so that a term is ± 1 if and only if it corresponds to a dimer tiling, and all other terms are 0. Choosing certain weights such that flips do not change the sign, it is observed in [FHNQ11] that the Hopf number invariant mod 2 is equal to the sign of the term in the Pfaffian (and one can check that the trit increments this number). From this they note that if M is defined analogously to in two dimensions, then in 3D

$$\text{Pf}(M) = A - B$$

where $A + B$ would be the partition function. The term A counts tilings with Hopf number $0 \pmod 2$ and B counts tilings with Hopf number $1 \pmod 2$.

In [KS22], the number $A - B$ is called the *defect*. The definition of the *twist invariant* discussed above is extended to dimensions $d > 3$ as the sign of the appropriate Kasteleyn determinant [KS22, Definition 3.1].

One can check by enumerating the equations for a single cube (i.e. 12 edges) that it is not possible to choose 12 nonzero weights so that the six flips (corresponding to σ with sign -1)

and four trits (corresponding to σ with sign $+1$) contained in the cube all preserve the sign of the term in the Pfaffian. In fact the six flip equations plus one trit equation have no simultaneous solution with all nonzero weights.

More generally, there is a complete characterization of which graphs admit *Pfaffian weights* and thereby make it possible to compute the partition function (which is a priori a permanent) as a determinant or Pfaffian of a re-weighted matrix. It is shown that a bipartite graph G admits Pfaffian weights if and only if it does not “contain” $K_{3,3}$ [Lit75]. Here “contain” means G can be modified (by replacing a collection of disjoint paths of edges containing an even number of vertices with a single edges) to a graph H which has $K_{3,3}$ as a subgraph. One can see that in this sense \mathbb{Z}^3 contains $K_{3,3}$, and hence does not have Pfaffian weights. The class of graphs that have Pfaffian weights can also be described in a way so that the Pfaffian is computable by a polynomial-time algorithm [RST99].

3.3 Loop shift Markov chain for uniform sampling

In two dimensions, uniformly random dimer tilings of finite simply connected regions can be efficiently simulated by a Markov chain that generates random flips, see [RT00]. As we have seen, dimer tilings of topologically trivial finite regions in dimensions $d > 2$ are not connected under flips, and it is an open question even for very simple regions whether flips and trits are sufficient. Here we describe a different, non-local Markov chain method to generate uniform random dimer tilings. The algorithm works in any dimension and for regions that are not simply connected, and is how the simulations in the introduction are generated. The simple move executed at each step of our chain is to construct a “random loop” in the given dimer tiling, and “shift” the tiles along the loop. This is a well-known construction in computer science, see for instance [Bro86, Section 3]. In the physics literature, see also [HKMS03] for Monte Carlo simulations of dimers in three dimensions based on algorithms from [KM03, DK95].

Given a dimer tiling τ of a finite region $R \subset \mathbb{Z}^3$, a *loop* γ in τ is a sequence of distinct edges $e_0, e_1, \dots, e_{k-1} \in \tau$ where the odd vertex of e_i is adjacent to the even vertex of e_{i+1} for all $i \in \mathbb{Z}/k\mathbb{Z}$ for some $k \geq 2$. A *loop shift* of τ along γ is a move which replaces edges along γ by their complementary edges. Specifically the resulting tiling is

$$\tau' = (\tau \setminus \{e_0, e_1, \dots, e_{k-1}\}) \cup \{f_0, f_1, \dots, f_{k-1}\}$$

where $\{e_0, e_1, \dots, e_{k-1}\} \cup \{f_0, f_1, \dots, f_{k-1}\}$ form a loop in \mathbb{Z}^3 . Since R is finite, given any two dimer tilings τ, σ of R the double dimer tiling (τ, σ) is a finite collection of double edges and loops $\gamma_1, \dots, \gamma_n$ of finite length. Loop shifting τ along γ_i for each of these transforms τ into σ . In particular, we have shown that

Proposition 3.3.1. *Let τ and σ be tilings of a finite set $R \subset \mathbb{Z}^3$. Then τ can be transformed into σ by a finite sequence of loop shifts.*

Loop shift Markov chain M . Given that any two tilings of a finite region $R \subset \mathbb{Z}^3$ differ by a finite sequence of loop shifts, we define a Markov chain M where one step proceeds as follows:

- Start with some dimer tiling τ of the region R .

- Sample an odd vertex in R uniformly at random. Start a path by following the tile from τ at this point.
- Uniformly at random choose a direction (other than the one we came from), and move in that direction for the next step.
- Repeat this (following the tile from τ , then following a uniform random choice, etc) until the path hits itself to form a loop. Call the loop γ .
- Drop any initial segment of the path which is not part of the loop γ . Then shift along γ , switching the tiles from τ for the random choices that we made along the path, and replace τ with σ which differs from τ only along γ .

By Proposition 3.3.1, M is an irreducible Markov chain and hence has a stationary distribution π . A bound on the mixing time of M is not known, see Problem 9.0.2.

Theorem 3.3.2. *The stationary distribution π of M is the uniform distribution on dimer tilings of R .*

Proof. Let P be its transition matrix. It is sufficient to prove that P is symmetric. If τ, σ are tilings such that $P(\tau, \sigma) \neq 0$, then they differ along a single loop γ .

Suppose that λ is a connected alternating-tile path in τ which consists of an initial segment α plus the loop γ . $P(\tau, \sigma)$ is a sum of the probability of paths λ of this form. We will show that the probability of generating λ in τ is the same as the probability of generating λ' in σ , where λ' has the same initial segment as λ , then traverses γ with the reverse orientation.

Let v_1, \dots, v_{2n} be the vertices along λ . Note that the vertices with odd index are odd, and out of these we follow a tile from τ . The vertices with even index are even, and out of these we follow a random choice. Thus the probability of generating the path λ in τ is $\prod_{k=1}^n \frac{1}{\deg(v_{2k})-1}$.

The sequence of vertices along λ' is the same, just in a different order. However the even vertices are still the sites where we make a random choice of direction to follow, so the probability of generating the path λ' in σ is also $\prod_{k=1}^n \frac{1}{\deg(v_{2k})-1}$.

Hence $P(\tau, \sigma) = P(\sigma, \tau)$. □

3.4 Local move connectedness on the torus and k-Gibbs measures

Here we discuss local move connectedness for dimer tilings of the torus, which is not simply connected. For any tiling τ of the d -dimensional torus \mathbb{T}^d , there is a standard, natural way to associate a homology class $[a(\tau)] \in H_1(\mathbb{T}^d)$. For each $i = 1, \dots, d$, let P_i be any plane with normal vector η_i , the i^{th} unit coordinate vector. Let \mathbb{T}_n^d denote the $n = n_1 \times n_2 \times \dots \times n_d$ torus in dimension d . Without loss of generality, n_1 is even. Let τ_0 be the tiling of \mathbb{T}_n^d where all tiles $t \in \tau_0$ are of the form $t = ((2x, y, z), (2x+1, y, z))$. With slight abuse of notation, we let $v_\tau(p) = v_\tau(e)e$ for the edge e incident to p containing a dimer (in particular $v_\tau(p)$ is one of the $2d$ unit coordinate vectors). For $i = 1, \dots, d$, we define

$$a_i(\tau) = \sum_{p \in P_i \cap \mathbb{T}_n^d} v_\tau(p) - v_{\tau_0}(p) = \sum_{p \in P_i \cap \mathbb{T}_n^d} v_\tau(p).$$

Since $v_\tau - v_{\tau_0}$ is divergence-free, this is independent of the choice of plane P_i normal to η_i . The second equality follows from the fact that v_{τ_0} contributes 0 to the overall sum. The homology class of τ is

$$[a(\tau)] := [a_1(\tau), \dots, a_d(\tau)] \in H_1(\mathbb{T}^d) \simeq \mathbb{Z}^d.$$

Note that a parallel pair of tiles contributes 0 total flow across any coordinate plane intersecting it. In particular, in any dimension $d > 1$, flips cannot change the homology class of a tiling of \mathbb{T}^d . However, when $d = 2$ the homology class is the only obstruction: if τ, τ' are tilings of an $n_1 \times n_2$ torus \mathbb{T}_{n_1, n_2}^2 and $[a(\tau)] = [a(\tau')]$, then τ, τ' are connected by a finite sequence of flips.

In dimension $d = 3$, the story is very different. In fact:

Proposition 3.4.1. *There is no finite collection of local moves that can connect all homologically equivalent dimer tilings of \mathbb{T}^3 .*

Remark 3.4.2. The authors of [FKMS22] exhibited a tiling of the $8 \times 8 \times 4$ torus with no flips or trits, obtained by stacking horizontal brickwork patterns of different orientations. We use similar stacked brickwork patterns (but with thicker layers) in our proof of Proposition 3.4.1.

Proof. The fundamental example is the following. Let τ be a tiling of $\mathbb{T}_{n_1, n_2, 4}^3$ where the first layer is an η_1 brickwork tiling, the second layer is an η_2 brickwork tiling, the third layer is a $-\eta_1$ brickwork tiling, and the fourth layer is a $-\eta_2$ brickwork tiling. By construction, $[a(\tau)] = (0, 0, 0)$. On the other hand, τ_0 also has $[a(\tau_0)] = (0, 0, 0)$, so τ and τ_0 are homologically equivalent. On the other hand, the length of the shortest alternating-tile loop in τ is $\min\{n_1, n_2, 4\}$. To see this, note that if the loop is homologically non-trivial, it must be long enough to visit at least three different horizontal layers. If it is homologically trivial, then its length must be at least $\min\{n_1, n_2, 4\}$.

More generally, for any $n = (n_1, n_2, 4n_3)$, we can construct a tiling τ of \mathbb{T}_n^3 which has n_3 layers of η_1 brickwork, followed by n_3 layers of η_2 brickwork, n_3 layers of $-\eta_1$ brickwork, and n_3 layers of $-\eta_2$ brickwork. Again $[a(\tau)] = (0, 0, 0)$, however the shortest contractible loop in τ has length $4n_3$ (since, again, it has to visit at least three different brickwork patterns). Therefore to connect τ, τ_0 we need loops of length at least $\min\{n_1, n_2, 4n_3\}$. These dimensions can be arbitrarily large, so this completes the proof. \square

By lifting this construction from \mathbb{T}^3 to \mathbb{R}^3 , we get the following corollary:

Corollary 3.4.3. *There is no finite collection of local moves which connects any two tilings of \mathbb{Z}^3 which differ at only finitely many places.*

Proof. Fix an integer $n > 0$. Tile all of \mathbb{Z}^3 with alternating brickwork layers so that there are n layers of η_1 brickwork, n layers of η_2 brickwork, n layers of $-\eta_1$ brickwork, and n layers of $-\eta_2$ brickwork. We denote the resulting tiling of \mathbb{Z}^3 by τ_n .

The shortest length of a cycle in τ_n is $4n$. Since there are finite cycles in τ_n , there exist tilings σ which differ from τ_n at only finitely many places. On the other hand, we need a local move of length at least $4n$ to make any change to τ_n . Since n is arbitrary this completes the proof. \square

Another interesting observation can be made from the example used in the proof of Proposition 3.4.1. A measure μ is *k-Gibbs* if for any box B with side length k , it holds that conditioned

on a tiling τ of $\mathbb{Z}^3 \setminus B$, μ is the uniform measure on tilings σ of B extending τ . If a measure is k -Gibbs for all k , then it is Gibbs.

In two dimensions, any two tilings of a $k \times k$ box (with the same boundary condition) are connected by some finite sequence of flips. Therefore if a measure on dimer tilings of \mathbb{Z}^2 is 2-Gibbs, then it is k -Gibbs for all k and hence Gibbs. The analogous statement does not hold in three dimensions.

Proposition 3.4.4. *For any integer $k \geq 2$ there exist k -Gibbs measures on Ω which are not Gibbs measures.*

Proof. Take $n = (n_1, n_2, n_3, n_4)$ and consider the tiling of \mathbb{Z}^3 which alternates between n_1 layers of η_1 bricks, n_2 layers of η_2 bricks, n_3 layers of $-\eta_1$ bricks, and n_4 layers of $-\eta_2$ bricks. Define a measure μ_m by averaging over translations by $\mathbb{Z}_{\text{even}}^3$ in the $m \times m \times m$ box and let μ be a subsequential limit as $m \rightarrow \infty$. The measure μ is invariant under the action of $\mathbb{Z}_{\text{even}}^3$. For $k \leq \min\{n_1, n_2, n_3, n_4\}$, μ is k -Gibbs since within any size k cube, a tiling sampled from μ is frozen for $k \leq \min\{n_1, n_2, n_3, n_4\}$. For $k > \min\{n_1, n_2, n_3, n_4\}$, μ still a.s. samples tilings which are brickwork patterns restricted to horizontal layers. However tilings of these larger boxes are not frozen, and are connected by shifting on finite loops to tilings which are not brickwork on every layer. Therefore μ is not k -Gibbs for $k > \min\{n_1, n_2, n_3, n_4\}$, hence μ is not Gibbs. \square

The construction in the proof works to construct a k -Gibbs-but-not-Gibbs measure for any mean current $s = (s_1, s_2, 0)$. A more complicated construction allows us to show that there exist k -Gibbs measures which are not Gibbs and correspond to an s in the interior of \mathcal{O} for which $s_1 s_2 s_3 \neq 0$. (Essentially one can arrange a periodic pattern of infinite non-intersecting taut paths like the ones shown in the introduction.) We have not found a construction that works for every $s \in \mathcal{O}$.

4 Measures with boundary mean current

Recall from Section 2.2 that $\mathbb{Z}_{\text{even}}^3$ -invariant measures on dimer tilings of \mathbb{Z}^3 come with a parameter called the *mean current*. This definition makes sense in any dimension d . When $d = 2$, the mean current is a 90-degree rotation of the height function *slope*, and in general it is valued in the convex polyhedron

$$\mathcal{O}_d = \{(s_1, \dots, s_d) : |s_1| + \dots + |s_d| \leq 1\}.$$

Recall that the mean current of a measure μ is defined in terms of tile densities (Definition 2.2.2). Given a standard basis $\eta_1, \eta_2, \dots, \eta_d$ of \mathbb{Z}^d denote by e_i the edge connecting 0 with η_i and $-e_i$ the edge connecting 0 with $-\eta_i$. The *mean current* of a measure $\mu \in \mathcal{P}(\Omega)$ is an element of \mathbb{R}^d such that its i^{th} -coordinate is

$$(s(\mu))_i = \mu(e_i \in \tau) - \mu(-e_i \in \tau).$$

If $s \in \partial\mathcal{O}_d$ we say that s is a *boundary mean current*. In terms of tiles, a measure μ has boundary mean current if and only if with probability 1 it samples at most one of the two possible tile types in each coordinate direction. The purpose of this section is to describe




Figure 18: Parts of five strips drawn on the dual graph (left) and as a tiling (right)

ergodic Gibbs measures with boundary mean current in dimension three. Using this, we compute the entropy function $\text{ent}(\cdot)$ in 3D restricted to $\partial\mathcal{O} = \partial\mathcal{O}_3$ (Theorem 4.2.7).

We will see that measures with boundary mean current in 2D and 3D are qualitatively very different. While the EGMs with boundary mean current in two dimensions all have zero entropy, EGMs with boundary mean current $s \in \partial\mathcal{O}$ in three dimensions can have positive entropy when s is contained in the interior of a face of $\partial\mathcal{O}$. Further, in three dimensions for any value a between 0 and $\text{ent}(s)$, there exists an EGM μ with specific entropy $h(\mu) = a$.

Despite these differences, in 2D and 3D the general principle is that measures with boundary mean current in dimension d correspond to sequences of measures on a $(d - 1)$ -dimensional lattice. This is easy to see in 2D, and we use it as a warm-up for the 3D version.

4.1 Review: EGMs with boundary mean current in two dimensions

Call the four possible tile directions in two dimensions (east, west) and (north, south). It is sufficient to describe measures with boundary mean current (s_1, s_2) for which $s_1, s_2 \geq 0$ and $s_1 + s_2 = 1$, i.e. measures that sample only north and east tiles. The first step is to understand what tilings containing only north and east tiles look like.

For an even point (x_1, x_2) , the north tile connects it to $(x_1, x_2 + 1)$ and the east tile connects it to $(x_1 + 1, x_2)$. In other words, north and east tiles always connect points along the line $x_1 + x_2 = 2c$ to points along the line $x_1 + x_2 = 2c + 1$. Therefore a tiling consisting of only north and east tiles can be partitioned into an infinite sequence of complete dimer tilings of *strips* $T_c = \{(x_1, x_2) : x_1 + x_2 = 2c \text{ or } 2c + 1\}$.

Along each strip, there are only two complete dimer tilings: one where the tiles are all east, and one where the tiles are all north. As such, any tiling τ with only north and east tiles consists of a sequence of choices of north or east tiles along the strips. See Figure 18.

All tilings τ of \mathbb{Z}^2 containing only north and east tiles are *frozen*, meaning they contain no finite cycles. To see this, note that if τ contains a finite cycle, then local move connectedness (see Section 3) implies it could be broken down into cycles of length 2. However a cycle of length 2 requires a north-south or east-west pair of tiles, which is not possible if the tiling contains only north and east tiles. Since tilings containing only north and east tiles are frozen, any measure μ which a.s. samples such tilings is automatically Gibbs. Three useful




Figure 19: (1) Cubes from a slab of \mathbb{Z}^3 visible from above the slab in a dimer tiling τ of \mathbb{Z}^3 , (2) tiles from τ drawn on the hexagonal lattice as edges colored pink, blue and orange, (3) the same tiles drawn as lozenges obtained by taking the Voronoi cells containing these edges, and (4) a key giving the translation between lozenge tiles and 3D dimer bricks.

observations follow from this discussion:

- All ergodic Gibbs measures with boundary mean current in two dimensions have zero entropy. In other words, entropy is zero when restricted to $\partial\mathcal{O}_2$.
- There is a bijection between 1) Gibbs measures on dimer tilings of \mathbb{Z}^2 that contain only E and N tiles and 2) measures on integer-indexed $\{N, E\}$ sequences. Any sample of a process taking value E with proportion s_1 and N with proportion s_2 , corresponds to a sample of a Gibbs measure on dimer tilings (obtained by placing N and E tiles on consecutive strips) of \mathbb{Z}^2 with mean current (s_1, s_2) and vice versa.
- There is also a bijection between 1) ergodic Gibbs measures on dimer tilings of \mathbb{Z}^2 that contain only E and N tiles and have mean current $(p, 1 - p)$ and 2) ergodic measures on integer-indexed $\{N, E\}$ sequences where the origin has probability p of being assigned to E .

4.2 EGMs with boundary mean current in three dimensions

Now we will consider the three dimensional case. Let the types of tiles be (east, west), (north, south), (up, down). Without loss of generality we consider measures with boundary mean current that almost surely sample only north, east, and up tiles, i.e. mean current $s = (s_1, s_2, s_3)$ with $s_1 + s_2 + s_3 = 1$, $s_1, s_2, s_3 \geq 0$.

For an even point (x_1, x_2, x_3) , an east tile connects it to $(x_1 + 1, x_2, x_3)$, a north tile connects it to $(x_1, x_2 + 1, x_3)$, and an up tile connects it to $(x_1, x_2, x_3 + 1)$. Therefore a tiling in 3D using only these three tile types corresponds to a sequence of tilings on two-dimensional *slabs*,

$$L_c = \{(x_1, x_2, x_3) : x_1 + x_2 + x_3 = 2c \text{ or } 2c + 1\}.$$

These slabs turn out to be a familiar two-dimensional lattice, namely the hexagonal lattice (with dimers viewed as edges) or the dual triangular lattice (with each dimer is viewed as a “lozenge” obtained as the union of two adjacent triangles), see Figure 19.

In the following, given a dimer tiling τ of a slab L_c , we will say that a particular tile type (north, east or up) has density s_i if the proportion of tiles of that type in $\tau \cap [-n, n]^3$ converges to s_i as $n \rightarrow \infty$. Similarly we can define the density for lozenge tilings.

Proposition 4.2.1. *For each $c \in \mathbb{Z}$, the slab L_c is a copy of the hexagonal lattice. There is a correspondence between tilings τ of \mathbb{Z}^3 which use only north, east and up tiles restricted to L_c and lozenge tilings. This correspondence takes a tiling of L_c with density (s_1, s_2, s_3) of the north, east and up tiles to a lozenge tiling where the density of the three lozenge tiles is also (s_1, s_2, s_3) .*

Remark 4.2.2. There is a completely analogous correspondence for $s \in \partial\mathcal{O}$ when some of the components of s are negative. If the signs of s are $(\epsilon_1, \epsilon_2, \epsilon_3)$ then a tiling with boundary mean current s would restrict to a lozenge tiling on $\{(x_1, x_2, x_3) : \epsilon_1 x_1 + \epsilon_2 x_2 + \epsilon_3 x_3 = 2c \text{ or } 2c + 1\}$. To simplify the presentation, some of the results in this section are stated for $s_1, s_2, s_3 \geq 0$, but the analogous statements hold for all $s \in \partial\mathcal{O}$.

Proof. Here we view the tiling as a collection of edges. Since τ uses only north, east and up tiles, the restriction $\tau_c = \tau|_{L_c}$ is a complete tiling of L_c . A single cube C in the \mathbb{Z}^3 lattice intersects four layers of the form $x_1 + x_2 + x_3 = a$. Let \mathcal{C}_a be the collection of cubes in \mathbb{Z}^3 which intersect the layers $x_1 + x_2 + x_3 = a - 1, a, a + 1, a + 2$. By construction, $L_c \subset \mathcal{C}_{2c}$.




Figure 20: Three adjacent cubes in \mathcal{C}_{2c} , with intersection with L_c in orange.

For each $C \in \mathcal{C}_{2c}$, $C \cap L_c$ is a hexagon, hence the faces of L_c are hexagons. By observation we see that any two adjacent hexagons meet in an edge, any three adjacent hexagons meet at a vertex, and there are no collections of > 3 adjacent hexagons. Hence L_c is a copy of the hexagonal lattice. Finally Figure 19 gives the correspondence between the north, east and up tiles with the three kinds of lozenges which preserves their densities. \square

Recall that \mathcal{P}^s denotes the set of $\mathbb{Z}_{\text{even}}^3$ -invariant probability measures on dimer tilings of \mathbb{Z}^3 of mean current $s = (s_1, s_2, s_3)$. We add subscripts G and e to denote Gibbs and ergodic measures respectively. Consider the group

$$\mathbb{Z}_{\text{loz}} = \mathbb{Z}_{\text{even}}^3 \cap \{(x_1, x_2, x_3) : x_1 + x_2 + x_3 = 0\}.$$

Let \mathcal{P}_{loz} denote the space of probability measures on dimer tilings of the slab L_0 (i.e. lozenge tilings) which are invariant under the \mathbb{Z}_{loz} action. The *slope* of a measure ρ on lozenge tilings is the triple $s = (s_1, s_2, s_3)$ of expected densities of the three types of lozenges with respect to ρ . A lozenge tiling slope satisfies $s_1, s_2, s_3 \geq 0$ and $s_1 + s_2 + s_3 = 1$.

We abuse notation slightly and write $s(\mu)$ to mean the mean current or slope depending on what space μ is a measure on. To reduce notation issues, for the rest of the subsection we denote measures on dimer tilings of \mathbb{Z}^3 by μ or ν and measures on lozenge tilings by ρ or λ .

In this section, we use the notation τ_B to mean τ restricted to $B \subset \mathbb{Z}^3$.

Proposition 4.2.3. *Suppose μ is an ergodic Gibbs measure on dimer tilings of \mathbb{Z}^3 with mean current $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$, $s_1, s_2, s_3 \geq 0$. Let ρ_0 be the marginal measure of μ on the slab L_0 . Then ρ_0 is a \mathbb{Z}_{loz} -invariant Gibbs measure with lozenge tiling slope $s(\rho_0) = s$.*

Remark 4.2.4. Since μ is $\mathbb{Z}_{\text{even}}^3$ -invariant, μ is invariant under the \mathbb{Z} -action of translating by $(0, 0, 2c)$, which takes L_0 to L_c . Therefore ρ_c and ρ_0 are identically distributed for all $c \in \mathbb{Z}$.

Proof. Since $\mathbb{Z}_{\text{loz}} \subset \mathbb{Z}_{\text{even}}^3$, ρ_0 is \mathbb{Z}_{loz} -invariant. Consider a finite connected set $B \subset L_0$ with boundary ∂B in \mathbb{Z}^3 . Suppose $\tau \in \Omega$ is a tiling in the support of μ , implying that it is a tiling using only north, east, and up tiles. Since μ is a Gibbs measure, and since there is μ -a.s. no tile in τ connecting L_0 and L_c for $c \neq 0$, we have for any tiling $\sigma \in \Omega$,

$$\mu(\sigma_B \mid \tau_{\mathbb{Z}^3 \setminus B}) = \mu(\sigma_B \mid \tau_{\partial B \cap L_0}) = \rho_0(\sigma_B \mid \tau_{\partial B \cap L_0}).$$

In the above we use the notation that for a tiling $\sigma \in \Omega$ and a set $A \subset \mathbb{Z}^3$, σ_A means σ restricted to A . Since μ is a Gibbs measure, the left hand side is uniform. Therefore ρ_0 is also a Gibbs measure.

Relating $s(\rho_0)$ to $s(\mu)$ is straightforward. Recall from Section 2.2 that $s_0(\tau)$ denotes the direction of the tile at the origin in τ , and that $s(\mu) = \mathbb{E}_\mu[s_0(\tau)]$. The same function s_0 can be used to compute the slope of a lozenge tiling measure, and $s(\rho_0) = \mathbb{E}_\rho[s_0(\sigma)]$ where σ is a full-plane lozenge tiling.

Let $\tau_0 = \tau_{L_0}$ be τ restricted to L_0 . Since μ has boundary mean current, τ_0 is a full-plane lozenge tiling μ a.s. and $s_0(\tau) = s_0(\tau_0)$. Thus

$$s(\mu) = \mathbb{E}_\mu[s_0(\tau)] = \mathbb{E}_{\rho_0}[s_0(\tau_0)] = s(\rho_0).$$

□

To show that $h(\mu) = h(\rho_0)$ (Proposition 4.2.5), we use the fact that any Gibbs measure can be uniquely decomposed into *extreme* Gibbs measures [Geo11, Theorem 7.26]. Extreme Gibbs measures are the extreme points of the convex set of Gibbs measures (analogous to how ergodic measures are the extreme points of the convex set of invariant measures). A Gibbs measure is extreme if and only if it is tail trivial [Geo11, Theorem 7.7].

If ρ is a Gibbs measure, there is a unique weight function g_ρ on the extreme Gibbs measures which gives its *extreme Gibbs decomposition*,

$$\rho = \int \lambda \, dg_\rho(\lambda).$$

This decomposition means that sampling from a Gibbs measure ρ can be thought of as a two step process: 1) sample an extreme Gibbs component λ from dg_ρ , 2) sample a tiling τ from λ . Given a tiling τ sampled from a Gibbs measure ρ , we can a.s. recover the extreme Gibbs component λ that τ was sampled from [Geo11, Theorem 7.12]. If λ is the extreme Gibbs component that τ is sampled from, we say that τ is *generic* for λ .

Proposition 4.2.5. *Suppose μ is an EGM on Ω with mean current $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$, $s_1, s_2, s_3 \geq 0$. Let ρ_c be the marginal measure of μ on the slab L_c for $c \in \mathbb{Z}$. Sampling a tiling τ from μ induces a choice of extreme Gibbs component λ_c of ρ_c for all $c \in \mathbb{Z}$. For each $c \in \mathbb{Z}$, let $\tau_c = \tau_{L_c}$.*

1. Conditional on the choice of extreme Gibbs component λ_c of ρ_c for each $c \in \mathbb{Z}$, the samples $(\tau_c)_{c \in \mathbb{Z}}$ are independent.
2. For any $c \in \mathbb{Z}$, $h(\mu) = h(\rho_c)$.

Remark 4.2.6. Since ρ_0 and ρ_c are identically distributed, $h(\rho_c) = h(\rho_0)$ for all $c \in \mathbb{Z}$. Thus it suffices to prove (2) for $c = 0$. We also note that we could have used the ergodic decomposition instead of the extreme Gibbs decomposition to prove this theorem. The upshot of using the extreme Gibbs decomposition is that conditional on a choice of extreme Gibbs component λ_c on each slab L_c , the samples τ_c from λ_c for all $c \in \mathbb{Z}$ are independent. Conditioned on a choice of ergodic component η_c on each slab L_c , only the samples from η_c with $s(\eta_c) = (l_1^c, l_2^c, l_3^c)$ and $l_1^c, l_2^c, l_3^c > 0$ are necessarily independent [She05, Theorem 9.1.1].

Proof of Proposition 4.2.5. Since μ is Gibbs, it has an extreme Gibbs decomposition

$$\mu = \int \nu \, dg_\mu(\nu).$$

Sampling a tiling τ from μ is equivalent to sampling an extreme Gibbs component ν of μ (from g_μ), and then sampling a tiling τ from ν . Since ν is tail-trivial, its marginal λ_c on L_c is also tail-trivial. For all $c \in \mathbb{Z}$, the extreme Gibbs decomposition of the marginal ρ_c can be written as

$$\rho_c = \int \lambda_c \, dg_{\rho_c}(\lambda_c)$$

where g_{ρ_c} is the extreme Gibbs decomposition of ρ_c .

Let $B_n = [-n, n]^3$. Since λ_0 is extreme Gibbs it is tail trivial, so

$$\lim_{m \rightarrow \infty} \mu(\tau_{B_n \cap L_0} \mid \tau_{\mathbb{Z}^3 \setminus (B_m \cap L_0)}, \lambda_0) = \lambda_0(\tau_{B_n \cap L_0}).$$

Therefore conditioned on λ_0 , $\tau_0 = \tau_{L_0}$ is independent of $\tau_{(\mathbb{Z}^3 \setminus L_0)}$. In particular, conditioned on the sequence of measures $(\lambda_c)_{c \in \mathbb{Z}}$ (equivalently, conditioned on choosing an extreme Gibbs component of μ), the tilings on each slab $(\tau_c)_{c \in \mathbb{Z}}$ are independent.

Now we relate the specific entropies. Recall from Section 2.3 that for a region $\Delta \subset \mathbb{Z}^3$ and an invariant measure μ on tilings of \mathbb{Z}^3 ,

$$H_\Delta(\mu) = - \sum_{\sigma \in \Omega(\Delta)} \mu(X(\sigma)) \log \mu(X(\sigma)),$$

where $\Omega(\Delta)$ is the free-boundary tilings of Δ , and $X(\sigma)$ is the collection of tilings of \mathbb{Z}^3 which extend σ . Taking $A_n(0) = B_n \cap L_0$, let $A_n(c) = A_n(0) + (0, 0, 2c)$, and finally let $A_{n,m} = \cup_{c=-m}^m A_n(c)$. It is well known that the specific entropy can be computed as

$$h(\mu) = \lim_{n \rightarrow \infty} |A_{n,n}|^{-1} H_{A_{n,n}}(\mu).$$

Instead of free-boundary tilings, we can choose $\tau \in \Omega$ and let $\Omega_\tau(\Delta) = \{\sigma \in \Omega(\Delta) : \sigma|_{\partial\Delta} = \tau\}$ be the tilings of Δ with boundary condition agreeing with τ . Then we define the entropy of μ given a fixed boundary condition τ :

$$H_\Delta(\mu|\tau) = - \sum_{\sigma \in \Omega_\tau(\Delta)} \mu(\sigma \mid \tau_{\mathbb{Z}^3 \setminus \Delta}) \log \mu(\sigma \mid \tau_{\mathbb{Z}^3 \setminus \Delta}).$$

Again for $A \subset \mathbb{Z}^3$, τ_A means τ restricted to A . We remark that this is not the usual definition of conditional entropy where we condition on a random variable or a sigma algebra. Instead we are fixing the value of the random variable, namely, the boundary condition of the tiling in Δ . Indeed, if τ is generic for an extreme Gibbs measure component ν of μ then we have that $H_\Delta(\mu|\tau) = H_\Delta(\nu|\tau)$. We will restrict this non-standard usage to this proof. It is standard that the specific entropy of μ can also be computed using this conditional definition as

$$h(\mu) = \lim_{n \rightarrow \infty} |A_{n,n}|^{-1} \left(\int_{\Omega} H_{A_{n,n}}(\mu|\tau) d\mu(\tau) + H_{\partial A_{n,n}}(\mu) \right) = \lim_{n \rightarrow \infty} |A_{n,n}|^{-1} \int_{\Omega} H_{A_{n,n}}(\mu|\tau) d\mu(\tau). \quad (10)$$

In the second equality, we use that the entropy term for $\partial A_{n,n}$ is of order n^2 so it does not contribute in the limit. Now we rewrite the argument of the limit using the extreme Gibbs decomposition.

$$\int_{\Omega} H_{A_{n,n}}(\mu|\tau) d\mu(\tau) = \int \int_{\Omega} H_{A_{n,n}}(\nu|\tau) d\nu(\tau) dg_{\mu}(\nu).$$

Recall that $\tau_c = \tau_{L_c}$. Conditional on sampling the process $(\lambda_c)_{c \in \mathbb{Z}}$ (equivalently, conditional on sampling ν), the samples $(\tau_c)_{c \in \mathbb{Z}}$ are independent. Thus

$$H_{A_{n,n}}(\nu|\tau) = \sum_{c=-n}^n H_{A_n(c)}(\lambda_c|\tau_c).$$

Recall that Ω_{loz} is the set of full-plane lozenge tilings. Thus

$$\int \int_{\Omega} H_{A_{n,n}}(\nu|\tau) d\nu(\tau) dg_{\mu}(\nu) = \sum_{c=-n}^n \int \int_{\Omega_{\text{loz}}} H_{A_n(c)}(\lambda_c|\tau_c) d\lambda_c(\tau_c) dg_{\rho_c}(\lambda_c) \quad (11)$$

$$= \sum_{c=-n}^n \int_{\Omega_{\text{loz}}} H_{A_n(c)}(\rho_c|\tau_c) d\rho_c(\tau_c). \quad (12)$$

Since ρ_c is equal in distribution to ρ_0 , for all $c \in \mathbb{Z}$,

$$|A_n(c)|^{-1} \int_{\Omega_{\text{loz}}} H_{A_n(c)}(\rho_c|\tau_c) d\rho_c(\tau_c) = |A_n(0)|^{-1} \int_{\Omega_{\text{loz}}} H_{A_n(0)}(\rho_0|\tau_0) d\rho_0(\tau_0).$$

At the same time,

$$\lim_{n \rightarrow \infty} |A_n(0)|^{-1} \int_{\Omega_{\text{loz}}} H_{A_n(0)}(\rho_0|\tau_0) d\rho_0(\tau_0) = h(\rho_0). \quad (13)$$

Therefore

$$\begin{aligned} h(\mu) &= \lim_{n \rightarrow \infty} |A_{n,n}|^{-1} \int_{\Omega} H_{A_{n,n}}(\mu|\tau) d\mu(\tau) \\ &= \lim_{n \rightarrow \infty} \frac{1}{2n+1} (2n+1) |A_n(0)|^{-1} \int_{\Omega_{\text{loz}}} H_{A_n(0)}(\rho_0|\tau_0) d\rho_0(\tau_0) = h(\rho_0). \end{aligned}$$

□

Recall from Section 2.3 that the *mean-current entropy function* $\text{ent} : \mathcal{O} \rightarrow [0, \infty)$ is defined by

$$\text{ent}(s) = \max_{\mu \in \mathcal{P}^s} h(\mu).$$

This entropy function plays a central role in our work and will be studied extensively in Section 7.

Let $\mathcal{T}_2 = \{s_1, s_2, s_3 \geq 0 : s_1 + s_2 + s_3 = 1\}$ be the space of possible lozenge tiling slopes. The *slope entropy function* $\text{ent}_{\text{loz}} : \mathcal{T}_2 \rightarrow [0, \infty)$ for lozenge tilings is defined by

$$\text{ent}_{\text{loz}}(s) = \max_{\rho \in \mathcal{P}_{\text{loz}}^s} h(\rho).$$

It was shown in [CKP01, Theorem 9.2] that ent_{loz} has the explicit form

$$\text{ent}_{\text{loz}}(s_1, s_2, s_3) = \frac{1}{\pi}(L(\pi s_1) + L(\pi s_2) + L(\pi s_3))$$

where $L : [0, \pi] \rightarrow \mathbb{R}$ is the Lobachevsky function given by

$$L(\theta) = - \int_0^\theta \ln(2 \sin(x)) dx.$$

Using this two dimensional result, we can explicitly compute ent on $\partial\mathcal{O}$. Let $\mathcal{E} \subset \partial\mathcal{O}$ denote the edges of \mathcal{O} .

Theorem 4.2.7. *For $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$,*

$$\text{ent}(s) = \text{ent}_{\text{loz}}(|s_1|, |s_2|, |s_3|) = \frac{1}{\pi}(L(\pi|s_1|) + L(\pi|s_2|) + L(\pi|s_3|)).$$

Further, if $s \notin \mathcal{E}$, then any measure μ realizing $h(\mu) = \text{ent}(s)$ is an ergodic Gibbs measure on Ω with respect to the $\mathbb{Z}_{\text{even}}^3$ action. If $s \in \mathcal{E}$, then $\text{ent}(s) = 0$.

It is well-known that ent_{loz} is strictly concave as a function of slope on the interior of allowed slopes [CKP01, Theorem 10.1]. Thus as an immediate corollary, we get that

Corollary 4.2.8. *Let \mathcal{F} be any face of $\partial\mathcal{O}$. The entropy function $\text{ent}(\cdot)$ is strictly concave on the interior of \mathcal{F} .*

Proof of Theorem 4.2.7. By Theorem 7.1.2, if $\mu \in \mathcal{P}^s$ satisfies $h(\mu) = \text{ent}(s)$, then μ is a Gibbs measure. While we include this result later in the paper for organizational reasons, it follows easily from the classical variational principle for Gibbs measures [LR69] (the only adaptation is that we are looking at the maximizer with a fixed mean current).

Without loss of generality assume that $s_1, s_2, s_3 \geq 0$. First suppose that $\mu \in \mathcal{P}^s$ is an EGM, and as usual let ρ_c denote its marginal on L_c . By Proposition 4.2.3 and Proposition 4.2.5,

$$s = s(\mu) = s(\rho_0) \quad \text{and} \quad h(\mu) = h(\rho_0).$$

Combining the results of [CKP01] and [She05],

- If s has $s_1, s_2, s_3 > 0$, then $\rho_0 \in \mathcal{P}_{\text{loz}}^s$ satisfies $h(\rho_0) = \text{ent}_{\text{loz}}(s)$ if and only if ρ_0 is the unique ergodic Gibbs measure of slope s , which we denote by λ_s .
- If s has $s_i = 0$ for some $i = 1, 2, 3$, then $h(\rho_0) = \text{ent}_{\text{loz}}(s) = 0$ for all $\rho_0 \in \mathcal{P}_{\text{loz}}^s$.

Therefore if $s_1, s_2, s_3 > 0$ and $h(\rho_0) = \text{ent}_{\text{loz}}(s)$, then by strict concavity of ent_{loz} [CKP01, Theorem 10.1], ρ_c are identically distributed and equal to λ_s a.s. By [She05, Theorem 9.1.1], when $s_1, s_2, s_3 > 0$, the unique ergodic Gibbs measure λ_s is an extreme Gibbs measure, and thus $\{\rho_c\}_{c \in \mathbb{Z}}$ is i.i.d. by Proposition 4.2.5. Alternatively if $s_i = 0$ for some i , then $h(\rho_0) = 0$, and hence $h(\mu) = 0$.

If μ is not ergodic with respect to the $\mathbb{Z}_{\text{even}}^3$ action, then it can be decomposed

$$\mu = \int_{\mathcal{P}_e} \nu \, dw_\mu(\nu),$$

where

$$s(\mu) = \int_{\mathcal{P}_e} s(\nu) \, dw_\mu(\nu).$$

Note that w_μ almost surely, $s(\nu)$ is contained in the same face of $\partial\mathcal{O}$ as $s = s(\mu)$. By the analysis above for an ergodic measure, if $s(\nu) \notin \mathcal{E}$ then $h(\nu) = \text{ent}(s(\nu)) = \text{ent}_{\text{loz}}(s(\nu))$ if and only if ν is an EGM of mean current $s(\nu)$ with marginals on each slab i.i.d. and equal to the lozenge tiling EGM of slope $s(\nu)$ (if $s(\nu) \in \mathcal{E}$, then $s(\nu) = 0$). Since ent_{loz} is strictly concave on the interior of allowed slopes, if s is contained in the interior of a face of $\partial\mathcal{O}$, then $h(\mu) = \text{ent}(s)$ if and only if μ is an ergodic Gibbs measure of mean current s . \square

As seen in the proof of Theorem 4.2.7, we get an explicit description of the entropy maximizers for $s \in \partial\mathcal{O}$. In contrast to two dimensions, the maximum entropy is positive for mean currents in the interior of faces.

Corollary 4.2.9. *Suppose $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$.*

- *If $s \in \mathcal{E}$ (i.e. $s_i = 0$ for some $i = 1, 2, 3$), then $h(\mu) = 0$ for any $\mu \in \mathcal{P}^s$.*
- *If $s_1 s_2 s_3 \neq 0$, then the entropy maximizer in \mathcal{P}^s is an ergodic Gibbs measure such that for all $c \in \mathbb{Z}$, $\rho_c = \lambda_s$ a.s., where λ_s is the unique ergodic Gibbs measure on lozenge tilings with slope $(|s_1|, |s_2|, |s_3|)$. Here ρ_c is the marginal on the slab $\{(x_1, x_2, x_3 : \epsilon_1 x_1 + \epsilon_2 x_2 + \epsilon_3 x_3 = 2c \text{ or } 2c + 1)\}$, where ϵ_i is the sign of s_i .*

It is also straightforward to show that there exist EGMs of a fixed boundary mean current with a range of different entropies.

Proposition 4.2.10. *Suppose $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$, $s_1, s_2, s_3 > 0$. Then for all $0 \leq \theta \leq 1$ there is an ergodic Gibbs measure μ such that $h(\mu) = \theta \text{ent}(s)$.*

Proof. Let $\rho_{\max}, \rho_1, \rho_2, \rho_3 \in \mathcal{P}_{\text{loz}}$ be EGMs on lozenge tilings of slopes (s_1, s_2, s_3) , $(1, 0, 0)$, $(0, 1, 0)$, $(0, 0, 1)$ respectively. Now consider an i.i.d. process $(\eta_c)_{c \in \mathbb{Z}}$ with state space

$$\{\rho_{\max}, \rho_1, \rho_2, \rho_3\}$$

such that the probability of ρ_{\max} is θ , and the probability of ρ_i is $(1 - \theta)s_i$ for $i = 1, 2, 3$.

Let μ be a measure on Ω given by taking a sample from $(\eta_c)_{c \in \mathbb{Z}}$, this gives a tiling of \mathbb{Z}^3 such that the restriction to each slab L_c is an independent sample from η_c . Clearly μ is a Gibbs measure on Ω . Since $(\eta_c)_{c \in \mathbb{Z}}$ is an i.i.d. process it is ergodic so μ is ergodic with respect to $\mathbb{Z}_{\text{even}}^3$. By Proposition 4.2.3 $s(\mu) = s$ and by Proposition 4.2.5 $h(\mu) = \theta \text{ent}(s)$. \square

We now summarize the results from this section to illustrate the similarities and differences with the two dimensional case.

- In three dimensions, EGMs of the same boundary mean current s can have different specific entropy values (Proposition 4.2.10).
- Every EGM μ on dimer tilings that contains only east, north, and up tiles gives rise to a Gibbs measure on integer-indexed stationary sequences of extreme Gibbs measures $(\lambda_c)_{c \in \mathbb{Z}}$ on lozenge tilings (Proposition 4.2.5).
- If $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$ is such that $s_1, s_2, s_3 \neq 0$ then the entropy-maximizing measure with mean current s is an EGM such that $(\lambda_c)_{c \in \mathbb{Z}}$ is an i.i.d. sequence of copies of the unique EGM on lozenge tilings with slope $(|s_1|, |s_2|, |s_3|)$ (Corollary 4.2.9).

5 Free-boundary tilings, asymptotic flows, and Wasserstein distance

This section sets up some of the key function-theoretic preliminaries for the large deviation principle in Section 8.

A domain is an open subset of \mathbb{R}^3 . Let $R \subset \mathbb{R}^3$ be a compact region which is the closure of a connected domain and has piecewise smooth boundary ∂R . We say that a grid region G is *scale n* if $G \subseteq \frac{1}{n}\mathbb{Z}^3$. If R_n is a scale n grid region, then with a slight abuse of notation we say that $R_n \supseteq R$ if the collection of $\frac{1}{n}$ -width cubes centered at points in R_n contains R . If τ is a tiling of R_n , we define the restriction of τ to R , denoted τ_R , to be the collection of tiles from τ which intersect R .

Definition 5.0.1. The *free-boundary tilings of R at scale n* are

$$T_n(R) := \bigcup_{R_n \supseteq R} \{\tau_R : \tau \text{ is a tiling of } R_n\}.$$

The *free-boundary tiling flows on R at scale n* are

$$TF_n(R) := \{f_\tau : \tau \in T_n(R)\}.$$

Finally, we define the space of all free-boundary tiling flows on R to be $TF(R) := \bigcup_{n \geq 1} TF_n(R)$.

The edges in $\frac{1}{n}\mathbb{Z}^3$ have length $\frac{1}{n}$. To ensure that the total flow of a tiling flow is roughly constant in n , we need the flow per edge of $f_\tau \in TF_n(R)$ to be of order $\frac{1}{n^3}$. We can achieve that by rescaling the flow by a factor of n^3 so that it has magnitude $\frac{5}{6n^3}$ on each matched edge and $\frac{1}{6n^3}$ on each unmatched edge.

Remark 5.0.2. Note that $TF_n(R)$ may contain elements that do not arise as restrictions of tilings of all of $\frac{1}{n}\mathbb{Z}^3$ to R . That is, there may be free-boundary tilings of R that cover R but do not extend to tilings of all of \mathbb{Z}^3 . (These might exist, for example, if R is a concave region.)

We define a metric on flows (Section 5.2), denoted d_W , which is an adaptation of *generalized Wasserstein distance* from signed measures to flows. Intuitively we want to consider two flows f, g “close” if we don’t have to change the flow of f too much—either by moving flow over, or by adding or deleting it—to transform it into g . This is what $d_W(f, g)$ will measure. In

terms of this metric, the main question of this section is: if $f_n \in TF_n(R)$ for all $n \in \mathbb{N}$, what are the possible limits of the form $\lim_{n \rightarrow \infty} f_n$?

We show (Theorem 5.3.1) that any fine-mesh limit of free-boundary tiling flows on R is an *asymptotic flow on R* , defined by:

Definition 5.0.3. Let R be a compact region which is the closure of a connected domain and has piecewise smooth boundary. We say that f is an *asymptotic flow on R* if it satisfies the following properties:

- f is a Borel-measurable vector field with support contained in R ;
- f is valued in \mathcal{O} (since f is measurable, this means that f is valued in \mathcal{O} Lebesgue-a.e.);
- f is divergence-free in the interior of R , i.e. $\operatorname{div} f = 0$ as a distribution (so for any smooth function ϕ compactly supported in the interior of R , $\int_R \phi \operatorname{div} f := \int_R \nabla \phi \cdot f = 0$.)

We denote the set of asymptotic flows on R by $AF(R)$.

In Theorem 5.3.4 we will show that $(AF(R), d_W)$ is a compact metric space. In Sections 5.4 and 5.5, we define a boundary value operator T (*trace operator*) which takes a flow to its boundary value on ∂R . In fact we do something more general, and define the trace of a flow for any compact, piecewise smooth surface contained in R . After defining T for asymptotic flows, we define the space of asymptotic flows with boundary value b , denoted $AF(R, b)$, and show that it is compact with respect to d_W (Corollary 5.4.9).

The boundary value operator is defined in different but analogous ways for asymptotic flows (Section 5.4) and tiling flows (Section 5.5). The main essential result about T is that these definitions are compatible and that T is uniformly continuous (Theorem 5.5.7).

We remark that the main important property of Wasserstein distance in our analysis is that it metrizes weak convergence, and that it thereby formalizes the intuitive notions that the scaling limits of tiling flows are asymptotic flows, and that boundary values depend continuously on the flow. While the Wasserstein metric and other transportation metrics have a number of additional special properties, we do not use this theory here. All the properties of the Wasserstein metric that we use are described in Section 5.1.

5.1 Background on (generalized) Wasserstein distance

The original *Wasserstein distance* or *earth-movers distance* is a metric on probability measures on a fixed metric space. Suppose that (X, d) is a compact, separable metric space. The L^1 Wasserstein distance W_1 is a metric on $\mathcal{P}(X)$, the space of probability measures on X and is given by

$$W_1(\mu, \nu) := \inf_{\gamma \in \Gamma(\mu, \nu)} \int_{X \times X} d(x, y) d\gamma(x, y)$$

where $\Gamma(\mu, \nu)$ is the collection of all couplings of μ and ν . Intuitively W_1 measures the cost—i.e. how much mass and how far it has to be moved—required to transform μ into ν by redistributing the mass of μ . This metric was developed in the theory of optimal transport and has been applied in many different contexts including probability, Riemannian geometry, and image processing. For more see [Vil09].

We will define and use versions of Wasserstein distance for flows and their boundary values. In the next section, we define a mapping between flows and measures, where flow (with direction) corresponds to mass (with sign). The measures corresponding to flows do not necessarily have the same mass and can be signed. Given this, our Wasserstein distance on flows will be based on a version of *generalized Wasserstein distance*.

Let $\mathcal{M}(\mathbb{R}^d)$ denote the space of Borel regular measures on \mathbb{R}^d with finite total mass. In [PR14] and [PR16], they define a generalized Wasserstein distance on $\mathcal{M}(\mathbb{R}^d)$ by introducing an L^1 cost for adding and deleting mass. It is denoted $W_1^{1,1}$ and defined as

$$W_1^{1,1}(\mu, \nu) = \inf_{\tilde{\mu}, \tilde{\nu}} |\mu - \tilde{\mu}| + |\nu - \tilde{\nu}| + W_1(\tilde{\mu}, \tilde{\nu})$$

where the infimum is taken over $\mathcal{M}(\mathbb{R}^d)$.

In [AMS11] the L^1 Wasserstein distance was generalized to signed probability measures. This metric is denoted \mathbb{W}_1 . If μ, ν are signed measures with Jordan decompositions $\mu = \mu_+ - \mu_-$ and $\nu = \nu_+ - \nu_-$, then

$$\mathbb{W}_1(\mu, \nu) = W_1(\mu_+ + \nu_-, \nu_+ + \mu_-).$$

In fact, note that this definition does not depend on the decomposition of the measures μ, ν .

In [PRT19], they combine these to give a definition of Wasserstein distance for signed measures of different masses. Let $\mathcal{M}^s(\mathbb{R}^d)$ denote the space of signed Radon measures on \mathbb{R}^d , i.e. measures μ that can be written $\mu_+ - \mu_-$ for $\mu_{\pm} \in \mathcal{M}(\mathbb{R}^d)$. Denoted $\mathbb{W}_1^{1,1}$, the generalized Wasserstein distance on $\mathcal{M}^s(\mathbb{R}^d)$ is defined

$$\mathbb{W}_1^{1,1}(\mu, \nu) = W_1^{1,1}(\mu_+ + \nu_-, \nu_+ + \mu_-).$$

This is the definition of Wasserstein distance that we will use in this paper. We note a few important facts about $\mathbb{W}_1^{1,1}$ that we will use.

Lemma 5.1.1 (See [PRT19, Lemma 18]). *If $\mu, \nu, \rho \in \mathcal{M}^s(\mathbb{R}^d)$, then*

$$\mathbb{W}_1^{1,1}(\mu, \nu) = \mathbb{W}_1^{1,1}(\mu + \rho, \nu + \rho).$$

Proposition 5.1.2 (See [PRT19, Proposition 23]). *Let*

$$\mathcal{C}_b^{0,Lip} = \{f : \mathbb{R}^d \rightarrow \mathbb{R} : f \text{ continuous, bounded, Lipschitz}\}.$$

Then

$$\mathbb{W}_1^{1,1}(\mu, \nu) = \sup \left\{ \int_{\mathbb{R}^d} \varphi d(\mu - \nu) : \varphi \in \mathcal{C}_b^{0,Lip}, \|\varphi\|_{\infty} \leq 1, \|\varphi\|_{Lip} \leq 1 \right\}.$$

From this it clearly follows that

Corollary 5.1.3. *If $\lim_{n \rightarrow \infty} \mathbb{W}_1^{1,1}(\mu_n, \mu) = 0$, then μ_n converges weakly to μ .*

The non-signed generalized Wasserstein distance $W_1^{1,1}$ metrizes weak convergence for tight sequences of measures in $\mathcal{M}(\mathbb{R}^d)$ [PR14, Theorem 13], as does the original L^1 Wasserstein distance for probability measures [Vil09, Theorem 6.9]. With signed measures, slightly stranger behavior can occur in general, see e.g. [PRT19, Remark 26]. However we show that the

Wasserstein distance for flows defined below does metrize weak convergence, see Remark 5.2.5.

For $R \subset \mathbb{R}^d$, we define $\mathcal{M}(R)$ to be the Radon measures supported in R , and $\mathcal{M}^s(R)$ to be signed Radon measures supported in R . We let $\mathcal{M}_{\text{ac}}(R)$ (resp. $\mathcal{M}_{\text{ac}}^s(R)$) denote the Radon measures (resp. signed Radon measures) supported in R and absolutely continuous with respect to Lebesgue measure on \mathbb{R}^d . We say that $\mathcal{M}_{\text{ac}}^s(R, a, b)$ denotes absolutely continuous signed measures with densities valued between a and b . By Lemma 5.1.1, for any $a < 0$ and $b > 0$,

$$(\mathcal{M}_{\text{ac}}^s(R, a, b), \mathbb{W}_1^{1,1}) \cong (\mathcal{M}_{\text{ac}}(R, 0, b-a), \mathbb{W}_1^{1,1}) \quad (14)$$

as metric spaces. This identification has some useful consequences. From Equation (14) and [PR14, Proposition 15] it follows that

Proposition 5.1.4. *If R is compact, then $(\mathcal{M}_{\text{ac}}^s(R, a, b), \mathbb{W}_1^{1,1})$ is a compact metric space for $a, b \in \mathbb{R}$.*

5.2 Wasserstein distance for flows

Let R be a compact region which is the closure of a connected domain and has piecewise smooth boundary. If $f_\tau \in TF_n(R)$, then f_τ is supported in $B_{2/n}(R) = \{x : d(x, R) \leq 2/n\}$. We will define a correspondence between

1. vector fields f on $R \subset \mathbb{R}^3$ or $f_\tau \in TF_n(R)$, and
2. triples of signed measures (μ_1, μ_2, μ_3) supported in $B_{2/n}(R)$.

The idea is that the flow of the vector field in coordinate direction i corresponds to mass of the i^{th} measure, with sign coming from the direction of the flow. We define Wasserstein distance on vector fields through this correspondence:

Definition 5.2.1. The *Wasserstein distance on flows*, denoted d_W , is the sum of the generalized Wasserstein distances between the component measures. For any two vector fields f, g with corresponding triples of measures (μ_1, μ_2, μ_3) and (ν_1, ν_2, ν_3) , we define

$$d_W(f, g) := \mathbb{W}_1^{1,1}(\mu_1, \nu_1) + \mathbb{W}_1^{1,1}(\mu_2, \nu_2) + \mathbb{W}_1^{1,1}(\mu_3, \nu_3).$$

To complete the definition of the metric, we need to define the correspondences between vector fields and triples of measures. There will be two definitions, one for a measurable vector field on R and one for a discrete vector field $f_\tau \in TF_n(R)$. Let $x = (x_1, x_2, x_3)$ denote a point in \mathbb{R}^3 .

Definition 5.2.2. (Measures corresponding to a measurable vector field.) Let f be a measurable vector field supported in R . The components of $f(x) = (f_1(x), f_2(x), f_3(x))$ are measurable functions, and we define the corresponding triple of measures (μ_1, μ_2, μ_3) by

$$d\mu_i(x) = f_i(x) dx_1 dx_2 dx_3 \quad i = 1, 2, 3$$

where $dx_1 dx_2 dx_3$ denotes Lebesgue measure on \mathbb{R}^3 .

Definition 5.2.3. (Measures corresponding to a free-boundary tiling flow on R .) Suppose that $f = f_\tau \in TF_n(R)$ for some n . Let η_1, η_2, η_3 be the positively-oriented unit basis vectors. Orient all the edges e of $\frac{1}{n}\mathbb{Z}^3$ to be parallel to η_i , which we denote by $e \parallel \eta_i$. (Recall that changing the orientation of e changes the sign of $f(e)$.) The triple of measures (μ_1, μ_2, μ_3) corresponding to f is given by

$$d\mu_i(x) = \sum_{e \parallel \eta_i} \frac{2}{n^2} f(e) \mathbb{1}_e(x) dx_i \quad i = 1, 2, 3$$

where $\mathbb{1}_e$ denotes the indicator of the edge $e \in \frac{1}{n}\mathbb{Z}^3$, and dx_i is 1-dimensional Lebesgue measure in the direction of η_i . Note that μ_i is supported in $B_{2/n}(R)$.

Remark 5.2.4. The scaling factor $\frac{2}{n^2}$ ensures that each edge e such that $e \parallel \eta_i$ contributes $\frac{2f(e)}{n^3}$ total mass to μ_i . The normalization is justified by looking at the extreme examples corresponding to the brickwork tilings (i.e. tilings where all tiles are the same type). Each cube in the $\frac{1}{n}\mathbb{Z}^3$ mesh can be viewed as corresponding to its lower left edge. In the brickwork pattern, exactly half of these cubes will have a dimer in the lower left edge.

Remark 5.2.5. Now that Wasserstein distance on flows is defined, we can explain why it metrizes weak convergence of the component measures. We do this by explaining how we could “shift” everything to have positive mass and use Equation (14). For asymptotic flows, we can just add a copy of the 3-dimensional Lebesgue measure $dx_1 dx_2 dx_3$. For tiling flows, we note that we could have defined the corresponding measures to be positive by translating the mean-current octahedron \mathcal{O} by $\eta_1 + \eta_2 + \eta_3$. After the translation, a scale n tiling flow measure would take values in $\{1/(3n^2), 5/(3n^2), 7/(3n^2), 11/(3n^2)\}$ instead of $\{-5/(3n^2), -1/(3n^2), 1/(3n^2), 5/(3n^2)\}$.

In terms of adding measures, translating \mathcal{O} corresponds to adding a copy of 1-dimensional Lebesgue $\frac{2}{n^2} dx_i$ along each edge $e \parallel \eta_i$ in $\frac{1}{n}\mathbb{Z}^3$ to the scale n tiling flow measure $d\mu_i$. In the scaling limit as $n \rightarrow \infty$, this sum of 1-dimensional Lebesgue measures converges in $\mathbb{W}_1^{1,1}$ to $dx_1 dx_2 dx_3$. By Equation (14), this implies the “translated” tiling flows measures converge (i.e. ones defined on the translated \mathcal{O}) to the “translated” asymptotic flow measures (i.e. ones shifted by adding $dx_1 dx_2 dx_3$) in $W_1^{1,1}$ if and only if the tiling flow measures converge to the asymptotic flow measures in $\mathbb{W}_1^{1,1}$. Since $W_1^{1,1}$ metrizes weak convergence [PR14, Proposition 15], d_W metrizes weak convergence of the component measures corresponding to tiling and asymptotic flows.

Proposition 5.2.6. *The measures corresponding to tiling flows are divergence-free on the interior of R in the sense of distributions, i.e. if f is a tiling flow with corresponding measures (μ_1, μ_2, μ_3) , then for any ϕ smooth and supported in a compact set C contained in the interior of R ,*

$$\int_R \frac{\partial \phi}{\partial x_1} d\mu_1 + \int_R \frac{\partial \phi}{\partial x_2} d\mu_2 + \int_R \frac{\partial \phi}{\partial x_3} d\mu_3 = 0.$$

Proof. For $i = 1, 2, 3$, let $e_1^i, \dots, e_{k_i}^i$ be the edges from $\frac{1}{n}\mathbb{Z}^3$ such that $e_j^i \parallel \eta_i$, is oriented parallel to η_i , and which intersect R . Let (a_j^i, b_j^i) be the endpoints of e_j^i such that $b_j^i - a_j^i = \eta_i$. By the fundamental theorem of calculus,

$$\sum_{i=1}^3 \int_R \frac{\partial \phi}{\partial x_i} d\mu_i = \frac{2}{n^2} \sum_{i=1}^3 \sum_{j=1}^{k_i} (\phi(b_j^i) - \phi(a_j^i)) f(e_j^i)$$

If $v = a_j^i$ or b_j^i is not contained in the interior of R , then $\phi(v) = 0$. Therefore we can rewrite the above as a sum over vertices $v \in \frac{1}{n}\mathbb{Z}^3$ contained in the interior of R :

$$\sum_{i=1}^3 \int_R \frac{\partial \phi}{\partial x_i} d\mu_i = \frac{2}{n^2} \sum_v \phi(v) F(v),$$

where $F(v)$ is a sum (with appropriate signs) of the six $f(e)$ terms for e incident to v . We show $F(v) = 0$.

Let e_i^-, e_i^+ denote the edges incident to v and oriented parallel to η_i . Let e_i^+ be the one for which the orientation parallel to η_i coincides with the orientation even to odd. Then

$$F(v) = \sum_{i=1}^3 f(e_i^+) - f(e_i^-).$$

But this is equal to $\sum_{\tilde{e} \ni v} f(\tilde{e})$, where the edges \tilde{e} incident to v are all oriented even to odd. Therefore $F(v) = 0$ since f is divergence-free as a discrete vector field, see Equation (5). \square

Next we prove a lemma about generalized Wasserstein distance for signed measures, in the case that both signed measures correspond to either tiling or asymptotic flows. This is an elementary result that we will use repeatedly.

Lemma 5.2.7. *Suppose that μ and ν are measures supported on a common compact set K corresponding to components of tiling or asymptotic flows. Suppose there is a partition of K into sets $\mathcal{B} = \{B_1, \dots, B_M\}$ of diameter at most ϵ such that $|\mu(B) - \nu(B)| < \delta$ for all $B \in \mathcal{B}$. If one of the measures corresponds to a scale n tiling flow, then we require that $\frac{1}{n} \leq \epsilon$. Then*

$$\mathbb{W}_1^{1,1}(\mu, \nu) \leq M(10\epsilon^4 + \delta).$$

Proof. Let $\mu = \mu_+ - \mu_-$ and $\nu = \nu_+ - \nu_-$ be decompositions into positive measures and recall that

$$\mathbb{W}_1^{1,1}(\mu, \nu) = W_1^{1,1}(\mu_+ + \nu_-, \mu_- + \nu_+).$$

Let $\tilde{\mu} = \mu_+ + \nu_-$ and $\tilde{\nu} = \mu_- + \nu_+$. To get an upper bound for the distance, it suffices to give a method for redistributing and deleting mass to transform $\tilde{\mu}$ into $\tilde{\nu}$.

We proceed as follows: transform $\tilde{\mu} |_{B_1}$ into $\tilde{\nu} |_{B_1}$, the cost of this is at most $W_1^{1,1}(\tilde{\mu} |_{B_1}, \tilde{\nu} |_{B_1})$. Denote the new version of $\tilde{\mu}$ by $\tilde{\mu}'$. $\tilde{\mu}'$ will agree with $\tilde{\mu}$ on $R \setminus B_1$ and with $\tilde{\nu}$ on B_1 . Next transform $\tilde{\mu}'$ into $\tilde{\nu}$ on B_2 . This will cost at most $W_1^{1,1}(\tilde{\mu}' |_{B_2}, \tilde{\nu} |_{B_2}) \leq W_1^{1,1}(\tilde{\mu} |_{B_2}, \tilde{\nu} |_{B_2})$ with equality if B_2 is disjoint from B_1 . Iterating this we get that

$$\mathbb{W}_1^{1,1}(\mu, \nu) \leq \sum_{j=1}^k W_1^{1,1}(\tilde{\mu} |_{B_j}, \tilde{\nu} |_{B_j})$$

Now we just have to compute the distance for a single B_j . First spend $\delta > 0$ to delete the difference in mass on B_j . The total mass of μ, ν on any $B \in \mathcal{B}$ is bounded by $10\epsilon^3$, and the furthest it would need to move is ϵ . Therefore

$$W_1^{1,1}(\tilde{\mu} |_{B_j}, \tilde{\nu} |_{B_j}) \leq 10\epsilon^4 + \delta.$$

Summing over j gives the result. \square

An inverse version of the bound in Lemma 5.2.7 also holds, but with a constant depending on the small region B .

Lemma 5.2.8. *Suppose $B \subset R$ is a connected region with piecewise smooth boundary. If μ, ν are component measures of tiling or asymptotic flows and $\mathbb{W}_1^{1,1}(\mu, \nu) < \delta$, then there is a constant $C(B)$ depending only on B such that*

$$\mathbb{W}_1^{1,1}(\mu|_B, \nu|_B) < \delta + (C(B) + 1)\delta^{1/2}.$$

Remark 5.2.9. The constant $C(B)$ is not hard to understand and control. The term $C(B)\delta^{1/2}$ is bounded by 2 times the volume of the $\delta^{1/2}$ annulus with inner boundary ∂B .

Proof. The redistribution, addition, and deletion of mass $\mu \rightarrow \nu$ gives a redistribution $\mu|_B \rightarrow \nu|_B$, except any mass moved into or out of B now becomes an L^1 cost rather than a cost proportional to distance moved. Let $f(r)$ be the amount of flow moved distance r into or out of B by the $\mu \rightarrow \nu$ redistribution. Then

$$\mathbb{W}_1^{1,1}(\mu|_B, \nu|_B) \leq \delta + \int_0^\infty f(r)dr.$$

On the other hand,

$$\int_0^\infty r f(r)dr < \delta.$$

We split the integral we want to bound into two pieces:

$$\int_0^\infty f(r)dr = \int_0^{\delta^{1/2}} f(r)dr + \int_{\delta^{1/2}}^\infty f(r)dr.$$

Since μ, ν are measures corresponding to components of tiling or asymptotic flows, we have

$$\int_0^{\delta^{1/2}} f(r)dr \leq C(B)\delta^{1/2}$$

(this quantity is proportional to the volume of the $\delta^{1/2}$ annulus around B) and

$$\delta^{1/2} \int_{\delta^{1/2}}^\infty f(r)dr \leq \int_{\delta^{1/2}}^\infty r f(r)dr < \delta.$$

Combining these gives the desired bound. \square

Lemma 5.2.10. *Let ν_n be a sequence of signed measures supported in R which converges in $\mathbb{W}_1^{1,1}$ to another measure ν . Further suppose the ν_n are absolutely continuous with respect to 3-dimensional Lebesgue measure, and their densities $g_n(x)$ take values in $[-m, M]$. Then ν is also absolutely continuous with respect to 3-dimensional Lebesgue measure.*

Proof. By Corollary 5.1.3, if $\mathbb{W}_1^{1,1}(\nu_n, \nu) \rightarrow 0$ as $n \rightarrow \infty$, then ν_n converges to ν in the weak topology.

Define an operator $Q : C^\infty(R) \rightarrow \mathbb{R}$ by integrating against ν :

$$Q(h) := \int h d\nu.$$

Using Cauchy-Schwarz and the fact that $d\nu_n = g_n(x) dx$ we have that

$$Q(h) = \lim_{n \rightarrow \infty} \int h g_n dx \leq \limsup_{n \rightarrow \infty} \|h\|_{L^2(R)} \|g_n\|_{L^2(R)} \leq \text{Vol}(R)^{1/2}(M + m) \|h\|_{L^2}$$

Therefore Q extends to an operator on $L^2(R)$. By the Riesz representation theorem this means there exists an L^2 function g such that $Q(h) = \langle h, g \rangle = \int h g dx$. Therefore $d\nu(x) = g(x)dx$, and ν is absolutely continuous with respect to Lebesgue measure on \mathbb{R}^3 . \square

Proposition 5.2.11. *Suppose that $f, g \in AF(R)$ are continuous and satisfy $|f(x) - g(x)| < \epsilon$ for all $x \in R$. Then $d_W(f, g) < \epsilon \text{vol}(R)$.*

Proof. The d_W -distance from f to g is bounded by adding and subtracting mass for each of the component functions. Since pointwise they differ by at most ϵ , $d_W(f, g) \leq \epsilon \text{vol}(R)$. \square

5.3 Main theorems

Here we prove two of the theorems mentioned at the beginning of the section. First we show that fine-mesh limits of tiling flows are asymptotic flows.

Theorem 5.3.1. *Let $R \subset \mathbb{R}^3$ be a compact region which is the closure of a connected domain and has piecewise-smooth boundary. Let $f_n \in TF_{m_n}(R)$ be a free-boundary tiling flow on R at scale m_n with m_n going to infinity with n . Any d_W -subsequential limit of tiling flows $f_* = \lim_{k \rightarrow \infty} f_{n_k}$ is in $AF(R)$.*

Let $\mu_k = (\mu_k^1, \mu_k^2, \mu_k^3)$ be the measures corresponding to f_{n_k} and let $\mu_* = (\mu_*^1, \mu_*^2, \mu_*^3)$ be the measures corresponding to f_* . The main idea of the proof is to smoothen the measures $(\mu_k^1, \mu_k^2, \mu_k^3)$ corresponding to the tiling flow in an especially nice way, then apply Lemma 5.2.10 to say that their limits $(\mu_*^1, \mu_*^2, \mu_*^3)$ are absolutely continuous with respect to 3-dimensional Lebesgue measure. This shows that f_* is a measurable flow on R . Using our well-chosen smoothings, we will show that f_* has the other properties that an asymptotic flow must have. To make the argument easier to digest, we break down the construction of the smoothing into two lemmas.

Lemma 5.3.2. *Let $S_N = [0, N - 1]^3$ (note that there are N vertices from \mathbb{Z}^3 on each edge of S_N). Let τ be a tiling of \mathbb{Z}^3 with tiling flow f_τ , and let (μ_1, μ_2, μ_3) be the corresponding measures. Then*

$$\frac{1}{\text{Vol}(S_N)} \left(\int_{S_N} d\mu_1, \int_{S_N} d\mu_2, \int_{S_N} d\mu_3 \right)$$

is valued in $(1 + O(1/N))\mathcal{O}$.

Proof. Let $E(S_N)$ denote the edges intersecting in S_N . All edges intersecting S_N in more than one point are contained in it and have length 1.

The measure μ_1 is supported on the edges parallel to η_1 , and thus

$$\int_{S_N} d\mu_1 = \sum_{E(S_N) \ni e \parallel \eta_1} 2f(e),$$

for edges e oriented parallel to η_1 . The results for μ_2, μ_3 are analogous.

For each $i = 1, 2, 3$, there are $N^2(N-1)$ edges from \mathbb{Z}^3 contained in S_N parallel to η_i . This number is always even, so half of these edges from $(N^2(N-1)/2)$ have even-to-odd orientation parallel to η_i and half have the opposite orientation. Let α_+ be the number of even-to-odd oriented tiles in $\tau \cap E(S_N)$ and α_- be the number of odd-to-even oriented tiles. Then

$$\int_{S_N} d\mu_1 = \sum_{E(S_N) \ni e \parallel \eta_i} 2f(e) = 2(\alpha_+ - \alpha_-).$$

Let s_i denote the fraction of tiles in the $+\eta_i$ direction in S_N minus the fraction of tiles in the $-\eta_i$ direction in S_N . Note that irrespective of the tiling we have that number of tiles in S_N is $N^2(N-1)/2 + O(N^2)$. Thus it follows that

$$s_1 = \frac{2(\alpha_+ - \alpha_-)}{N^2(N-1) + O(N^2)} = \frac{1}{N^2(N-1) + O(N^2)} \int_{S_N} d\mu_1.$$

A similar equation holds for s_2, s_3 . We have that $(s_1, s_2, s_3) \in \mathcal{O}$. Thus

$$\begin{aligned} \frac{1}{\text{Vol}(S_N)} \left(\int_{S_N} d\mu_1, \int_{S_N} d\mu_2, \int_{S_N} d\mu_3 \right) &= \frac{N^2(N-1) + O(N^2)}{(N-1)^3} (s_1, s_2, s_3) \\ &= \left(1 + \frac{O(N^2)}{(N-1)^3}\right) (s_1, s_2, s_3), \end{aligned}$$

for some constant C which is in $(1 + O(1/N))\mathcal{O}$. \square

If we smoothen the measures corresponding to the tiling flow over a partition consisting of boxes of the form in Lemma 5.3.2, we can construct smoothings that satisfy a very nice list of properties.

Lemma 5.3.3. *Let (μ_1, μ_2, μ_3) be measures corresponding to a tiling flow $f \in TF_n(R)$. For any $\epsilon > 0$, there exists $\nu = (\nu_1, \nu_2, \nu_3)$ satisfying the following properties:*

1. ν_i is supported in R for all $i = 1, 2, 3$;
2. ν_i has a density $g_i(x)$ with respect to Lebesgue measure on \mathbb{R}^3 for $i = 1, 2, 3$;
3. $g = (g_1, g_2, g_3)$ is valued in $(1 + O(\epsilon))\mathcal{O}$ as a distribution;
4. $d_W(\mu, \nu) < C(R)(\epsilon n)^{-1}$ where $C(R)$ is a constant depending only on R .

Proof. Choose N such $N = \lfloor \frac{1}{\epsilon} \rfloor$ and a partition of cubes $\mathcal{B} = \{B_1, \dots, B_M\}$ that cover R , where each B_i is an $N \times N \times N$ cube in $\frac{1}{n}\mathbb{Z}^3$ (we define the flow f to be 0 outside R). This can be done so that $M \sim n^3/N^3$. For $i = 1, 2, 3$ and all $B \in \mathcal{B}$, define

$$C_B^i := \frac{1}{\text{Vol}(B)} \int_B d\mu_i = \frac{(N-1)^3}{n^3} \int_B d\mu_i.$$

Define the densities of ν_i by

$$g_i(x) = C_B^i \quad \forall x \in B \cap R.$$

This satisfies properties 1 and 2. By Lemma 5.3.2, (g_1, g_2, g_3) is valued in $(1 + O(1/N))\mathcal{O} = (1 + O(\epsilon))\mathcal{O}$ as a distribution which completes property 3. Finally by Lemma 5.2.7 applied to the partition \mathcal{B} , we have that

$$d_W(\mu, \nu) \leq M(N/n)^4 \leq C(R)(n/N)^3(N/n)^4 \leq C(R) \frac{1}{\epsilon n}$$

where $C(R)$ is a constant depending only on R . \square

We now return to the proof of the theorem.

Proof of Theorem 5.3.1. Fix $\epsilon > 0$. Recall that $\mu_k = (\mu_k^1, \mu_k^2, \mu_k^3)$ is the triple of measures corresponding to f_{n_k} . Choose $\epsilon_k = n_k^{-1/2}$ and let $\nu_k = (\nu_k^1, \nu_k^2, \nu_k^3)$ be the measures constructed in Lemma 5.3.3 for $\epsilon_k > 0$. For k large enough so that $d_W(\mu_*, \mu_k) < \epsilon$, by the triangle inequality

$$d_W(\mu_*, \nu_k) \leq d_W(\mu_*, \mu_k) + d_W(\mu_k, \nu_k) \leq \epsilon + C(R)n_k^{-1/2}.$$

Therefore the triple of measures ν_k also converges to μ_* in d_W . By Lemma 5.2.10, there are functions f_*^i such that $\mu_*^i = f_*^i(x)dx$ for $i = 1, 2, 3$, so f_* is a measurable vector field. It remains for us to check the additional properties to show that f_* is an asymptotic flow.

Since the ν_k are supported in R for all k , so is f_* . By Lemma 5.3.3, the densities of ν_k are valued in $(1 + O(\epsilon_k))\mathcal{O}$, so the densities of μ_* are valued in \mathcal{O} (any open neighborhood is a continuity set, so we get that the averages of (f_*^1, f_*^2, f_*^3) are valued in $(1 + O(\epsilon_k))\mathcal{O}$ for all ϵ_k . This plus the Lebesgue differentiation theorem imply that f_* is valued in \mathcal{O}). On the other hand, convergence in d_W implies weak convergence of the component measures (Corollary 5.1.3). Since μ_k is divergence-free in the sense of distributions on the interior of R (Proposition 5.2.6) for all k , μ_* is also divergence-free in the sense of distributions on the interior of R . Thus $f_* \in AF(R)$. \square

Theorem 5.3.4. *The metric space $(AF(R), d_W)$ is compact.*

Proof. By Proposition 5.1.4, the space of triples of measures absolutely continuous with respect to Lebesgue measure, supported in R , and valued in \mathcal{O} is compact. Since $(AF(R), d_W)$ is a subspace of this, it suffices to show that it is closed. Suppose that $\mu_n = (\mu_n^1, \mu_n^2, \mu_n^3)$ is a sequence in $AF(R)$ with d_W -limit $\mu_* = (\mu_*^1, \mu_*^2, \mu_*^3)$. By Lemma 5.2.10, μ_*^i has a density $g_*^i(x)$ for each $i = 1, 2, 3$. Since convergence in $\mathbb{W}_1^{1,1}$ implies weak convergence (Corollary 5.1.3), $g_* = (g_*^1, g_*^2, g_*^3)$ is divergence-free. To show that g_* is valued in \mathcal{O} , note that any open ball $B \subset R$ is a continuity set, so since $\frac{1}{\text{Vol}(B)}(\int_B d\mu_n^1, \int_B d\mu_n^2, \int_B d\mu_n^3) \in \mathcal{O}$, the average of g_* over B is also valued in \mathcal{O} . Thus g_* is valued in \mathcal{O} by the Lebesgue differentiation theorem. Thus $(AF(R), d_W)$ is compact. \square

Now we will prove that any asymptotic flow can be approximated by a smooth flow which is divergence-free on a slightly smaller region. This is a standard construction which we provide for completeness as it will be used in the next subsection. Essentially all we need to do is to convolve the asymptotic flow with an appropriate smooth bump function. For this, given a region R and $\epsilon > 0$ we define

$$R_\epsilon := \{x \in R : d(x, \partial R) \geq \epsilon\}. \quad (15)$$

We will denote the smooth asymptotic flows on a region R by $AF^\infty(R) \subset AF(R)$. Given $f \in AF(R), g \in AF(R_\epsilon)$, we say that $d_W(f, g)$ is the distance between f and g , where g is extended to be 0 on $R \setminus R_\epsilon$.

Proposition 5.3.5. *Fix $f \in AF(R)$. For all $\epsilon > 0$ small enough, there is a smooth asymptotic flow $g \in AF^\infty(R_\epsilon)$ such that*

$$d_W(f, g) < K\sqrt{\epsilon}$$

where K is a constant depending only on R .

Proof. Consider a bump function $\phi \in C^\infty(B_\epsilon(0))$, that is, it is a non-negative smooth function such that $\phi|_{\partial B_\epsilon(0)} = 0$ and $\int_{B_\epsilon(0)} \phi(x) dx = 1$.

Let $g = f * \phi|_{R_\epsilon}$ and suppose that ψ is a smooth test function with compact support in the interior of R_ϵ . To check that g is divergence-free in the interior of R_ϵ we look at the integral

$$\begin{aligned} \int_{R_\epsilon} (\nabla\psi \cdot g)(x) dx &= \int_{R_\epsilon} \int_{B_\epsilon(0)} (\nabla\psi \cdot f)(x-y)\phi(y) dy dx \\ &= \int_{B_\epsilon(0)} \int_{R_\epsilon} (\nabla\psi \cdot f)(x-y)\phi(y) dx dy = 0. \end{aligned}$$

Here the last equality uses that f is divergence-free in the interior of R . We have that $g(x) = f * \phi(x)$ is an average of elements in \mathcal{O} . Since \mathcal{O} is convex it follows that g takes values in \mathcal{O} . Finally we estimate $d_W(f, g)$. The amount of mass which we might have to delete or add from $R \setminus R_\epsilon$ is bounded by $6(\text{Vol}(R \setminus R_\epsilon)) \leq c\epsilon^2$ where c depends only on ∂R . Now let (μ_1, μ_2, μ_3) and (ν_1, ν_2, ν_3) be the measures corresponding to f and g respectively. Let $\delta > 0$. We have that if B is a box of side length δ contained in R_ϵ then

$$\nu_i(B) = \int_B \int_{B_\epsilon(0)} f_i(x-y)\phi(y) dy dx = \int_{B_\epsilon(0)} \mu_i(B-y)\phi(y) dy.$$

It follows that $|\nu_i(B) - \mu_i(B)|$ is less than the volume of the annular region around B of radius ϵ . Thus we have that

$$|\nu_i(B) - \mu_i(B)| < C\epsilon\delta^2$$

where C is independent of ϵ and δ . Partition R_ϵ into boxes B_1, B_2, \dots, B_M ; where $M \sim \delta^{-3}$. By Lemma 5.2.7 we have that

$$\mathbb{W}_1^{1,1}(\mu_i, \nu_i) < M(10\delta^4 + C\epsilon\delta^2) + c\epsilon^2 \sim \delta^{-3}(10\delta^4 + C\epsilon\delta^2) + c\epsilon^2$$

Since δ is a free parameter, we can take $\delta \sim \sqrt{\epsilon}$ to complete the proof. \square

5.4 Boundary values of asymptotic flows

In this section we define the boundary values of asymptotic flows on ∂R . In fact we do something slightly more general, and define the restriction of an asymptotic flow (or tiling flow in the next subsection) on a whole class of surfaces, namely

$$\mathbb{S}(R) = \{\text{compact piecewise smooth surfaces contained in } R\}$$

Note that R is closed, so $\partial R \in \mathbb{S}(R)$. This general set up will make things easier to prove. We also use the trace operator for other surfaces in Section 8.6. Recall that $AF^\infty(R) \subset AF(R)$ is the smooth asymptotic flows.

Definition 5.4.1. We define the *trace operator* on smooth asymptotic flows

$$T : AF^\infty(R) \times \mathbb{S}(R) \rightarrow \mathcal{M}^s(R).$$

by

$$T(f, S)(x) = \langle f(x), \xi_S(x) \rangle d\sigma_S(x), \quad x \in S,$$

where $d\sigma_S$ denotes surface area measure on S and $\xi_S(x)$ denotes the L^2 unit normal vector to S at x .

We show that $T(\cdot, S)$ extends to a uniformly continuous map $(AF(R), d_W) \rightarrow (\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$ for all $S \in \mathbb{S}(R)$. We do this in three main steps:

- Show that $T(\cdot, S)$ is uniformly continuous on $AF^\infty(R)$ when S is a small patch.
- Extend this result to $AF(R)$ by approximation and compatibility results. Since we don't know if $AF^\infty(R)$ is dense in $AF(R)$, this requires slightly more care.
- Extend the uniform continuity result to general $S \in \mathbb{S}(R)$ by putting together the patches.

Proposition 5.4.2. *Suppose $S \in \mathbb{S}(R)$ and is such that there exists a nonzero vector v and a parameter $\theta > 0$ for which $S + tv$ is contained in R and disjoint for all $0 \leq t \leq \theta$. Then for all $\epsilon > 0$ there exists $\delta > 0$ such that for all $f, g \in AF^\infty(R)$ with $d_W(f, g) < \delta$, we have that*

$$\mathbb{W}_1^{1,1}(T(f, S), T(g, S)) < \epsilon.$$

Proof. Fix two parameters $\gamma_1, \gamma_2 > 0$ which we will specify at the end of the proof. Partition S into patches $\alpha_1, \dots, \alpha_M$ such that

- α_i is a smooth surface with piecewise smooth boundary for all $i = 1, \dots, M$.
- α_i has diameter at most γ_1 for all i , and $M \leq C\gamma_1^{-2}$ for some constant C depending on S .
- Let $\alpha_i(t) := \alpha_i + tv$ for $0 \leq t \leq \theta$. For all $i = 1, \dots, M$, $\alpha_i(t) \cap \alpha_i(s) = \emptyset$ for $s \neq t$.

Let $\mu_f = T(f, S)$, $\mu_g = T(g, S)$ and define $\Delta > 0$ by

$$\Delta := \sup_{1 \leq i \leq M} \left| \mu_f(\alpha_i) - \mu_g(\alpha_i) \right|. \quad (16)$$

By the two-dimensional version of Lemma 5.2.7,

$$\mathbb{W}_1^{1,1}(T(f, S), T(g, S)) \leq M(10\gamma_1^3 + \Delta) \leq 10C\gamma_1 + C\gamma_1^{-2}\Delta. \quad (17)$$

Note that the power of γ_1 is 3 instead of 4 because S is two-dimensional. It remains to find a bound for Δ in terms of $d_W(f, g)$.

If $h \in AF^\infty(R)$, then h is divergence-free and hence its flux through any closed surface is zero. This implies that there exists a threshold $\theta_{\gamma_2} \in (0, 1)$ such that for all $0 \leq t \leq \theta_{\gamma_2}$ and $h \in AF^\infty(R)$,

$$\sup_{1 \leq i \leq M} \left| \int_{\alpha_i} \langle h(x), \xi(x) \rangle d\sigma_{\alpha_i}(x) - \int_{\alpha_i(t)} \langle h(x), \xi(x) \rangle d\sigma_{\alpha_i(t)}(x) \right| \leq \gamma_2. \quad (18)$$

Here $\xi(x)$ is the normal vector on the surfaces $\alpha_i(t)$ with appropriate orientation and we are applying the divergence theorem to the boundary of the region $U_i = \cup_{s=0}^t \alpha_i(s)$. Since h takes values in the compact set \mathcal{O} , the threshold θ_{γ_2} can be taken independent of h . Applying this to $f, g \in AF^\infty(R)$ it follows that

$$\sup_{1 \leq i \leq M} \left| \int_{\alpha_i} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha_i}(x) - \int_{\alpha_i(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha_i(t)}(x) \right| \leq 2\gamma_2$$

for all $0 \leq t \leq \theta_{\gamma_2}$. Observe that the first term in the inequality is $\mu_f(\alpha_i) - \mu_g(\alpha_i)$. Integrating over $t \in (0, \theta_{\gamma_2})$,

$$\left| \int_0^{\theta_{\gamma_2}} \left(\int_{\alpha_i(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha_i(t)}(x) \right) dt - \theta_{\gamma_2}(\mu_f(\alpha_i) - \mu_g(\alpha_i)) \right| \leq 2\theta_{\gamma_2}\gamma_2.$$

Taking the supremum over $i = 1, \dots, M$ we get that

$$\Delta \leq 2\gamma_2 + \frac{1}{\theta_{\gamma_2}} \sup_{1 \leq i \leq M} \left| \int_0^{\theta_{\gamma_2}} \left(\int_{\alpha_i(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha_i(t)}(x) \right) dt \right|.$$

Suppose the supremum on the right hand side of the equation is achieved by the index i , and let $\alpha(t) := \alpha_i(t)$ to simplify notation. Then plugging this into (17) gives

$$\mathbb{W}_1^{1,1}(T(f, S), T(g, S)) \tag{19}$$

$$\leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2 + C\gamma_1^{-2}\theta_{\gamma_2}^{-1} \left| \int_0^{\theta_{\gamma_2}} \left(\int_{\alpha(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha(t)}(x) \right) dt \right|.$$

Since α_i is smooth, by an appropriate change of variables we can rewrite the integral above as an integral over $U = \cup_{t=0}^{\theta_{\gamma_2}} \alpha(t)$.

$$\begin{aligned} \int_0^{\theta_{\gamma_2}} \left(\int_{\alpha(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha(t)}(x) \right) dt &= \int_U \langle f(x) - g(x), \xi(x) \rangle \varphi(x) dx_1 dx_2 dx_3 \\ &= \sum_{k=1}^3 \int_U (f_k(x) - g_k(x)) \xi_k(x) \varphi(x) dx_1 dx_2 dx_3. \end{aligned}$$

Here $\varphi(x)$ is the factor coming from the Jacobian in the change of variables. Since $\alpha(t)$ is smooth for all $t \in [0, \theta_{\gamma_2}]$, $\varphi(x)$ and $\xi(x)$ are both smooth functions on U , so $\psi_k(x) := \xi_k(x)\varphi(x)$ is a smooth and therefore Lipschitz function on U . Let λ denote the maximum of Lipschitz constants of ψ_k , $k = 1, 2, 3$. Then by the dual definition of the Wasserstein metric (Proposition 5.1.2),

$$\int_0^{\theta_{\gamma_2}} \left(\int_{\alpha(t)} \langle f(x) - g(x), \xi(x) \rangle d\sigma_{\alpha(t)}(x) \right) dt \leq \lambda d_W(f|_U, g|_U).$$

By Lemma 5.2.8, there is a constant $C(U)$ such that if $d_W(f, g) < \delta$ then $d_W(f|_U, g|_U) < \delta + C(U)\delta^{1/2}$. Therefore substituting this in to Equation (19) gives that if $d_W(f, g) < \delta$ then

$$\mathbb{W}_1^{1,1}(T(f, S), T(g, S)) \leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2 + C\gamma_1^{-2}\theta_{\gamma_2}^{-1}\lambda(\delta + C(U)\delta^{1/2}).$$

Taking $\gamma_2 = \gamma_1^3$ completes the proof. \square

We will now prove that perturbing S by a small amount does not change the trace very much.

Proposition 5.4.3. *Let $S \in \mathbb{S}(R)$ be a surface satisfying the conditions of Proposition 5.4.2 for some vector v and threshold $\theta > 0$. For all $\epsilon > 0$ there exists $\delta > 0$ such that for all $t < \delta$ and $f \in AF^\infty(R)$, we have that*

$$\mathbb{W}_1^{1,1}(T(f, S), T(f, S + tv)) < \epsilon.$$

Proof. This proof is much simpler than that of Proposition 5.4.2. As in that proof, we take parameters $\gamma_1, \gamma_2 > 0$ to be fixed later and choose a partition of $\alpha_1, \dots, \alpha_M$ of S such that α_i has diameter less than γ_1 for all i , and $M \leq C\gamma_1^{-2}$ for some $C > 0$ independent of γ_1 . In addition, by choosing a larger constant C if necessary we can assume that

$$\sigma_S(\alpha_i) \leq C\gamma_1^2.$$

As in (18), there exists a threshold $\theta_{\gamma_2} \in (0, 1)$ such that for all $0 \leq t \leq \theta_{\gamma_2}$ and $f \in AF^\infty(R)$,

$$\sup_{1 \leq i \leq M} \left| T(f, S)(\alpha_i) - T(f, S + tv)(\alpha_i(t)) \right| \leq \gamma_2. \quad (20)$$

To get an upper bound for the distance between the traces on each patch, we give a method for redistributing, adding, and deleting mass to transform $T(f, S)|_{\alpha_i} = T(f, \alpha_i)$ into $T(f, \alpha_i(t))$. Note that these are both signed measures absolutely continuous with respect to σ_S and σ_{S+tv} respectively, and both have densities bounded between -1 and 1 . We can transform one to the other by (1) adding γ_2 flow, (2) moving flow distance at most $t + \gamma_1$. There is at most $4\sigma_S(\alpha_i)$ total flow from both measures. Hence

$$\mathbb{W}_1^{1,1}(T(f, S)|_{\alpha_i}, T(f, S + tv)|_{\alpha_i(t)}) \leq \gamma_2 + 4(t + \gamma_1)(\sigma_S(\alpha_i)) \leq \gamma_2 + 4C(t + \gamma_1)\gamma_1^2.$$

As in Lemma 5.2.7, by triangle inequality

$$\mathbb{W}_1^{1,1}(T(f, S), T(f, S + tv)) \leq M(\gamma_2 + 4C(t + \gamma_1)\gamma_1^2) \leq C\gamma_1^{-2}\gamma_2 + 4C^2(t + \gamma_1).$$

As before, by setting $\gamma_2 = \gamma_1^3$, the result follows. \square

Using Proposition 5.4.2 and Proposition 5.4.3, when $S \in \mathbb{S}(R)$ satisfies the conditions of Proposition 5.4.2 for a vector v and threshold $\theta > 0$, we can extend the definition of $T(S, \cdot)$ to all of $AF(R)$ as follows. Given $t > 0$, there exists K such that for all $k \geq K$, $S(t) \subset R_{1/k}$. Here recall that

$$R_{1/k} = \{x \in R : d(x, \partial R) \geq 1/k\}.$$

By Proposition 5.3.5, for any $f \in AF(R)$ we can find a sequence $g_k \in AF^\infty(R_{1/k})$ such that $d_W(f, g_k) \rightarrow 0$ as $k \rightarrow \infty$. For $t > 0$, we define

$$T(f, S(t)) := \lim_{k \rightarrow \infty} T(g_k, S(t)) \quad (21)$$

where the limit is taken with respect to the metric d_W on flows supported in R . Note that if $R' \subset R$, then the projection map $AF^\infty(R) \rightarrow AF^\infty(R')$ given by $f \mapsto f|_{R'}$ is continuous. By Proposition 5.4.2, T is uniformly continuous on $AF^\infty(R_{1/k})$ for any $k > 0$, so the limit in Equation (21) is independent of the approximating sequence $(g_k)_{k \geq 1}$ and converges to $T(f, S + tv)$ if $f \in AF^\infty(R)$. Since $(AF(R), d_W)$ is compact (Theorem 5.3.4), $T(\cdot, S(t)) : AF(R) \rightarrow \mathcal{M}^s(R)$ is a uniformly continuous map for $t > 0$.

Further, by a variant of Lemma 5.2.10 where 3-dimensional Lebesgue measure is replaced by the surface area measure $\sigma_{S(t)}$, for any $f \in AF(R)$, $T(f, S(t))$ is a signed measure absolutely continuous to $\sigma_{S(t)}$ with density bounded between -1 and 1 . By Proposition 5.4.3, for any $\epsilon > 0$, there exists $\delta > 0$ so that if $t, s > 0$ and $|t - s| < \delta$ then

$$\mathbb{W}_1^{1,1}(T(f, S(s)), T(f, S(t))) < \epsilon. \quad (22)$$

Thus we take another limit in the d_W topology to define

$$T(f, S) := \lim_{t \rightarrow 0} T(f, S(t)).$$

Since $AF(R)$ is compact and the extension above is continuous, we get analogs of Proposition 5.4.2 and Proposition 5.4.3 for $AF(R)$.

Proposition 5.4.4. *If $S \in \mathbb{S}(R)$ is a surface satisfying the conditions of Proposition 5.4.2 for a vector v and threshold $\theta > 0$, then*

1. *Using the extension described above,*

$$T(\cdot, S) : (AF(R), d_W) \rightarrow (\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$$

is uniformly continuous.

2. *Given $\epsilon > 0$ there exists $\delta > 0$ such that for all $t < \delta$ and $f \in AF(R)$,*

$$\mathbb{W}_1^{1,1}(T(f, S), T(f, S(t))) < \epsilon.$$

3. *For any $f \in AF(R)$, $T(f, S)$ is a signed measure absolutely continuous with respect to the surface area measure σ_S with density bounded between -1 and 1 .*

We now prove a compatibility result for traces on overlapping surfaces. Using this, we can extend the continuity theorems to the trace operator for any $S \in \mathbb{S}(R)$ by cutting S in patches which satisfy the conditions of Proposition 5.4.2.

Lemma 5.4.5. *Suppose that $S, S' \in \mathbb{S}(R)$ are surfaces satisfying the conditions of Proposition 5.4.2 such that $S' \subset S$. Then for any $f \in AF(R)$, then*

$$T(f, S) |_{S'} = T(f, S').$$

Proof. If $f \in AF^\infty(R)$, then the result follows immediately from the form given in Definition 5.4.1. For general $f \in AF(R)$ this follows from Proposition 5.4.4. \square

Finally we put together the pieces to prove uniform continuity of $T(\cdot, S)$ on $AF(R)$ for any $S \in \mathbb{S}(R)$.

Proposition 5.4.6. *For all $S \in \mathbb{S}(R)$, $T(\cdot, S)$ extends to a uniformly continuous map from $(AF(R), d_W)$ to $(\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$. Further, for all $f \in AF(R)$, $T(f, S)$ is a signed measure on S absolutely continuous to the surface measure σ_S with density bounded between -1 and 1 .*

Remark 5.4.7. In particular this holds for $S = \partial R$.

Proof. Fix $S \in \mathbb{S}(R)$. We can cover S with finitely many open surfaces S_1, \dots, S_k which all satisfy the conditions of Proposition 5.4.2. By Lemma 5.4.5, if $S_i \cap S_j \neq \emptyset$, then for all $f \in AF(R)$,

$$T(f, S_i) |_{S_i \cap S_j} = T(f, S_j) |_{S_i \cap S_j}$$

Therefore we can define the trace operator $T(S, \cdot)$ for any $S \in \mathbb{S}(R)$ by

$$T(f, S) |_{S_i} = T(f, S_i).$$

By Proposition 5.4.4, $T(f, S)$ is uniformly continuous as a function of $f \in AF(R)$ and has the desired form. \square

With this machinery, we can define the space of asymptotic flows with a fixed boundary value.

Definition 5.4.8. We say that $b \in \mathcal{M}^s(R)$ is a *boundary asymptotic flow* if $b \in T(AF(R), \partial R)$. Further, we define $AF(R, b)$ to be the space of asymptotic flows on R with boundary value b , i.e. $f \in AF(R)$ such that $T(f, \partial R) = b$.

As a corollary of Proposition 5.4.6 and Theorem 5.3.4 we get the following.

Corollary 5.4.9. *The metric space $(AF(R, b), d_W)$ is compact.*

5.5 Boundary values of tiling flows

Next we define the trace operator on tiling flows, and show that it is compatible with the definition for asymptotic flows. Suppose $f \in TF_n(R)$ and that $S \in \mathbb{S}(R)$ is a surface which intersects the lattice $\frac{1}{n}\mathbb{Z}^3$ transversely, i.e. S does not contain any vertices of $\frac{1}{n}\mathbb{Z}^3$ (if S contains a vertex, we translate the lattice slightly so that it does not). As usual let e denote an edge from $\frac{1}{n}\mathbb{Z}^3$ oriented from even to odd, and let $\xi(x)$ denote the normal vector to S at x .

Definition 5.5.1. If $f \in TF_n(R)$ and $S \in \mathbb{S}(R)$ is a surface intersecting $\frac{1}{n}\mathbb{Z}^3$ transversely, we define

$$T(f, S) = \sum_e \frac{2 \operatorname{sign}\langle \xi(x), e \rangle}{n^2} f(e) \delta(e \cap S).$$

Note that since S is transverse to $\frac{1}{n}\mathbb{Z}^3$, if $e \cap S$ is nonempty it is a single point.

Using Definition 5.5.1 for tiling flows and Definition 5.4.1 for asymptotic flows, the final goal of this section is to show that

$$T(\cdot, S) : (AF(R) \cup TF(R), d_W) \rightarrow (\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$$

is uniformly continuous for any $S \in \mathbb{S}(R)$ (see Theorem 5.5.7). The sequence of results in this section building up to this mirrors the sequence of results in the previous section. The discrete setting makes things slightly more complicated. The main new step is that we start by proving a result for the trace on planes, and extend to more general surfaces by approximating them with planes. Throughout, we assume that any surface S we consider intersects $\frac{1}{n}\mathbb{Z}^3$ transversely. Any time it does not, the trace is defined by perturbing the lattice slightly so that it does and then using Definition 5.5.1.

Proposition 5.5.2. *Suppose $P \in \mathbb{S}(R)$ is a compact piece of a plane with normal vector ξ , and there exists a threshold $\theta > 0$ such that $P(t) = P + t\xi$ is contained in R for all $t \in [0, \theta]$. Let $f_n \in TF_n(R)$ be a sequence of tiling flows such that $d_W(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$ for some $f \in AF^\infty(R)$. Then*

$$\lim_{n \rightarrow \infty} \mathbb{W}_1^{1,1}(T(f_n, P), T(f, P)) = 0.$$

Remark 5.5.3. The conditions here could be rephrased as saying that $P \in \mathbb{S}(R)$ is contained in a plane and satisfies the conditions of Proposition 5.4.2 with $v = \xi$.

Proof. As in Proposition 5.4.2, fix two parameters $\gamma_1, \gamma_2 > 0$ and partition P into patches $\alpha_1, \dots, \alpha_M$ such that

- α_i is a smooth surface with piecewise smooth boundary for all $i = 1, \dots, M$;
- α_i has diameter at most γ_1 for all $i = 1, \dots, M$, and $M \leq C\gamma_1^{-2}$ for some constant C depending on P ;
- Let $\alpha_i(t) := \alpha_i + t\xi$. For all $i = 1, \dots, M$, $\alpha_i(t) \cap \alpha_i(s) = \emptyset$ for $s \neq t$.

We define

$$\Delta_n = \sup_{1 \leq i \leq M} \left| T(f_n, P)(\alpha_i) - T(f, P)(\alpha_i) \right|.$$

By the two-dimensional version of Lemma 5.2.7,

$$\mathbb{W}_1^{1,1}(T(f_n, P), T(f, P)) \leq 10C\gamma_1 + C\gamma_1^{-2}\Delta_n.$$

Let $U_i(s) = \cup_{t=0}^s \alpha_i(t)$ be the paralleloiped region between $\alpha_i = \alpha_i(0)$ and $\alpha_i(s)$. Given γ_2 , we can find θ_{γ_2} small enough so that the number of edges from $\frac{1}{n}\mathbb{Z}^3$ hitting $\partial U_i(t) \setminus (\alpha_i(t) \cup \alpha_i)$ for any i is less than $\gamma_2 n^2 + K'n$, with constant K' depending on the length of $\partial \alpha_i$. Since the magnitude of f_n is of order $1/n^2$, for $t < \theta_{\gamma_2}$ the flow of f_n through $\partial U_i(t) \setminus (\alpha_i(t) \cup \alpha_i)$ is bounded by $O(n^{-1}) + \gamma_2$. Since any $f_n \in TF_n(R)$ is divergence-free as a discrete vector field, there exists a constant $K > 0$ so that for any f_n and $t \in (0, \theta_{\gamma_2})$,

$$\sup_{1 \leq i \leq M} \left| T(f_n, \alpha_i)(\alpha_i) - T(f_n, \alpha_i(t))(\alpha_i(t)) \right| \leq \gamma_2 + K/n. \quad (23)$$

Possibly choosing a smaller θ_{γ_2} , the same result holds for f without the K/n in the upper bound. By the triangle inequality,

$$\sup_{1 \leq i \leq M} \left| T(f_n, \alpha_i)(\alpha_i) - T(f, \alpha_i)(\alpha_i) - T(f_n, \alpha_i(t))(\alpha_i(t)) + T(f, \alpha_i(t))(\alpha_i(t)) \right| \leq 2\gamma_2 + K/n.$$

As in Proposition 5.4.2, integrating over $t \in (0, \theta_{\gamma_2})$ and solving for Δ_n gives

$$\Delta_n \leq 2\gamma_2 + K/n + \frac{1}{\theta_{\gamma_2}} \sup_{1 \leq i \leq M} \left| \int_0^{\theta_{\gamma_2}} (T(f_n, \alpha_i(t))(\alpha_i(t)) - T(f, \alpha_i(t))(\alpha_i(t))) dt \right|.$$

Let α_i be the patch where the supremum is achieved, and let $\alpha := \alpha_i$ to simplify notation. Then

$$\begin{aligned} \mathbb{W}_1^{1,1}(T(f_n, P), T(f, P)) &\leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2 + C\gamma_1^{-2}Kn^{-1} \\ &\quad + \frac{C\gamma_1^{-2}}{\theta_{\gamma_2}} \left| \int_0^{\theta_{\gamma_2}} T(f_n, \alpha(t))(\alpha(t)) - T(f, \alpha(t))(\alpha(t)) dt \right|. \end{aligned}$$

Finally we bound the integral in the last term. Let $U = U_i(\theta_{\gamma_2})$ to simplify notation. Recall that $\alpha(t) \subset P + t\xi$ is a piece of a plane, and has constant unit normal vector ξ . By Definition 5.4.1, $T(f, \alpha(t))(x) = \langle f(x), \xi \rangle$. Therefore letting $\sigma_{\alpha(t)}$ denote the surface area measure on $\alpha(t)$, and applying change of variables,

$$\begin{aligned} \int_0^{\theta} T(f, \alpha(t))(\alpha(t)) dt &= \int_0^{\theta} \int_{\alpha(t)} \langle f(x), \xi \rangle d\sigma_{\alpha(t)}(x) dt = \int_U \langle f(x), \xi \rangle dx_1 dx_2 dx_3 \\ &= \sum_{j=1}^3 \xi_j \mu_j(U), \end{aligned}$$

where $\xi = (\xi_1, \xi_2, \xi_3)$, $f = (f_1, f_2, f_3)$ and (μ_1, μ_2, μ_3) is the triple of measures corresponding to f . On the other hand, for the tiling flow f_n ,

$$\int_0^\theta T(f_n, \alpha(t))(\alpha(t)) dt = \sum_{j=1}^3 \xi_j \mu_j^n(U),$$

where $(\mu_1^n, \mu_2^n, \mu_3^n)$ is the triple of measures corresponding to f_n . Therefore

$$\begin{aligned} & \mathbb{W}_1^{1,1}(T(f_n, P), T(f, P)) \\ & \leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2 + C\gamma_1^{-2}Kn^{-1} + C\gamma_1^{-2}\theta_{\gamma_2}^{-1} \sum_{j=1}^3 |\xi_j| |\mu_j^n(U) - \mu_j(U)|. \end{aligned}$$

By Lemma 5.2.8, $d_W(f_n, f) \rightarrow 0$ implies that $|\mu_j^n(U) - \mu_j(U)| \rightarrow 0$ as $n \rightarrow \infty$ for $j = 1, 2, 3$. Taking $n \rightarrow \infty$ gives

$$\limsup_{n \rightarrow \infty} \mathbb{W}_1^{1,1}(T(f_n, P), T(f, P)) \leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2.$$

Setting $\gamma_2 = \gamma_1^3$ and taking $\gamma_1 \rightarrow 0$ completes the proof. \square

Next we prove a version of Proposition 5.4.3 for tiling flows and small patch surfaces as in Proposition 5.4.2.

Proposition 5.5.4. *Suppose that $S \in \mathbb{S}(R)$ satisfies the conditions of Proposition 5.4.2. For all $\epsilon > 0$ there exists $\delta > 0$ and $N > 0$ such that for all $t < \delta$, all $n \geq N$, and all $f \in TF_n(R)$,*

$$\mathbb{W}_1^{1,1}(T(f, S), T(f, S(t))) < \epsilon.$$

Proof. The proof is analogous to the proof of Proposition 5.4.3. Again we take parameters $\gamma_1, \gamma_2 > 0$ to be fixed later and a partition $\alpha_1, \dots, \alpha_M$ of S into patches of diameter at most γ_1 for all i , and such that $M \leq C\gamma_1^{-2}$ for some constant C independent of γ_1 . Analogous to Equation (23), given γ_2 we can find a threshold $\theta_{\gamma_2} > 0$ such that for all $0 \leq t \leq \theta_{\gamma_2}$ and all $f \in TF_n(R)$,

$$\sup_{1 \leq i \leq M} \left| T(f, S)(\alpha_i) - T(f, S(t))(\alpha_i(t)) \right| \leq \gamma_2 + K/n. \quad (24)$$

Using this, we get an upper bound for the distance by giving a method for redistributing, adding, and deleting mass to transform $T(f, S) |_{\alpha_i}$ into $T(f, S(t)) |_{\alpha_i(t)}$. Both measures are a sum of delta masses of weights $2/n^2(\pm 5/6)$ or $2/n^2(\pm 1/6)$. The number of delta masses in α_i or $\alpha_i(t)$ is bounded above by $\text{area}(\alpha_i)n^2$. Since α_i has diameter bounded by γ_1 , there is a constant $A > 0$ independent of α_i such that $\text{area}(\alpha_i) \leq A\gamma_1^2$. Hence the total mass in each patch is bounded between $-2A\gamma_1^2$ and $2A\gamma_1^2$. Hence adding $\gamma_2 + K/n$ mass plus moving at most $8A\gamma_1^2$ mass distance at most $t + \gamma_1$, we get the bound

$$\mathbb{W}_1^{1,1}(T(f, S) |_{\alpha_i}, T(f, S(t)) |_{\alpha_i(t)}) \leq \gamma_2 + K/n + 8A\gamma_1^2(t + \gamma_1)$$

Summing over i we get that

$$\begin{aligned} \mathbb{W}_1^{1,1}(T(f, S), T(f, S(t))) & \leq M(\gamma_2 + K/n + 8\gamma_1^2(t + \gamma_1)) \\ & = C\gamma_1^{-2}\gamma_2 + CK\gamma_1^{-2}n^{-1} + 8A(t + \gamma_1). \end{aligned}$$

Take $\gamma_2 = \gamma_1^3$ and γ_1, t small enough so that

$$(C + 8A)\gamma_1 + 8At < \epsilon/2. \quad (25)$$

Then take n large enough so that

$$CK\gamma_1^{-2}n^{-1} < \epsilon/2,$$

and the result follows with $\delta = \min\{\theta_{\gamma_1^3}, \frac{1}{8A}(\epsilon/2 - (C + 8A)\gamma_1)\}$. Here the first term in the minimum comes from (24) and the second comes from (25). \square

By approximation we can extend Proposition 5.5.2 to any surface $S \in \mathcal{S}(\mathbb{R})$. For technical reasons, we first show this for S contained strictly in the interior of R . Note that we also remove the condition that the limiting flow is smooth.

Proposition 5.5.5. *Suppose that $S \in \mathcal{S}(R)$ is contained strictly in the interior of R . Let $f_n \in TF_n(R)$ be a sequence of tiling flows such that $d_W(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$ for some $f \in AF(R)$. Then*

$$\lim_{n \rightarrow \infty} \mathbb{W}_1^{1,1}(T(f_n, S), T(f, S)) = 0.$$

Proof. As usual, let $\gamma_1, \gamma_2 > 0$ be small parameters to be fixed later. Since S is contained strictly in the interior of R , we can cover S by very small patch surfaces $\alpha_1, \dots, \alpha_M$ so that:

- Each α_i is smooth with piecewise smooth boundary,
- The diameter of α_i is at most γ_1 .
- There is a constant $C > 0$ such that $M \leq C\gamma_1^{-2}$.
- For all i , α_i satisfies the conditions of Proposition 5.4.2 for some threshold $\theta > 0$ with vector $v = \xi(q_i)$, where $\xi(q_i)$ is the normal vector to the surface at q_i for some $q_i \in \alpha_i$ with the property that the distance between q_i and any other $x \in \alpha_i$ is at most γ_1 .
- Let P_i denote the tangent plane to α_i at q_i , and let $\pi_i \subset P_i$ be the patch of the plane corresponding to projecting α_i onto P_i . Potentially making $\theta > 0$ or γ_1 smaller, we can assume that π_i also satisfies the conditions of Proposition 5.4.2 for $v = \xi(q_i)$. This is where we are using the fact that S is contained in the interior of R .

As usual we denote $\alpha_i(t) = \alpha_i + t\xi(q_i)$ and $\pi_i(t) = \pi_i + t\xi(q_i)$. Note by Definition 5.5.1 and Lemma 5.4.5 that

$$T(f_n, S) |_{\alpha_i} = T(f_n, \alpha_i) \quad \text{and} \quad T(f, S) |_{\alpha_i} = T(f, \alpha_i) \quad i = 1, \dots, M.$$

Define

$$\Delta_n := \sup_{1 \leq i \leq M} \left| T(f_n, S)(\alpha_i) - T(f, S)(\alpha_i) \right|.$$

By the two-dimensional version of Lemma 5.2.7,

$$\mathbb{W}_1^{1,1}(T(f_n, S), T(f, S)) \leq M(10\gamma_1^3 + \Delta_n) \leq 10C\gamma_1 + C\gamma_1^{-2}\Delta_n. \quad (26)$$

Since $f_n \in TF_n(R)$ is discrete divergence-free and S is compact and piecewise smooth, given γ_2 there exists $K > 0$ depending on S and a threshold θ_{γ_2} such that for all $0 \leq t \leq \theta_{\gamma_2}$,

$$\sup_{1 \leq i \leq M} \left| T(f_n, \alpha_i)(\alpha_i) - T(f_n, \alpha_i(t))(\alpha_i(t)) \right| \leq \gamma_2 + K/n. \quad (27)$$

By Proposition 5.4.4, up to making θ_{γ_2} smaller, for all $0 \leq t \leq \theta_{\gamma_2}$,

$$\sup_{1 \leq i \leq M} \left| T(f, \alpha_i)(\alpha_i) - T(f, \alpha_i(t))(\alpha_i(t)) \right| \leq \gamma_2. \quad (28)$$

Combining Equations (27) and (28), and as in Proposition 5.4.2 integrating over $t \in (0, \theta_{\gamma_2})$ then solving for Δ_n gives

$$\Delta_n \leq 2\gamma_2 + K/n + \frac{1}{\theta_{\gamma_2}} \sup_{1 \leq i \leq M} \left| \int_0^{\theta_{\gamma_2}} T(f_n, \alpha_i(t))(\alpha_i(t)) - T(f, \alpha_i(t))(\alpha_i(t)) dt \right|.$$

Let i be the index where the supremum is achieved, and let $\alpha(t) := \alpha_i(t)$ and $\pi(t) := \pi_i(t)$. We now bound

$$T(f_n, \alpha(t))(\alpha(t)) - T(f, \alpha(t))(\alpha(t))$$

using four terms. Let $(g_m) \in AF^\infty(R_{\epsilon_m})$, $\epsilon_m \rightarrow 0$ as $m \rightarrow \infty$, be a sequence of smooth asymptotic flows such that $\lim_{m \rightarrow \infty} d_W(g_m, f) = 0$. Since $\alpha(t), \pi(t)$ are contained strictly in the interior of R for $0 \leq t \leq \theta$, we can assume they are all contained in R_ϵ for some $\epsilon > 0$. In particular, for m large enough $\epsilon_m < \epsilon$ and hence g_m is defined on $\pi(t)$. We have for any $0 \leq t \leq \theta_{\gamma_2}$,

$$\begin{aligned} T(f_n, \alpha(t))(\alpha(t)) - T(f, \alpha(t))(\alpha(t)) &= T(f_n, \alpha(t))(\alpha(t)) - T(f_n, \pi(t))(\pi(t)) \\ &\quad + T(f_n, \pi(t))(\pi(t)) - T(g_m, \pi(t))(\pi(t)) \\ &\quad + T(g_m, \pi(t))(\pi(t)) - T(f, \pi(t))(\pi(t)) \\ &\quad + T(f, \pi(t))(\pi(t)) - T(f, \alpha(t))(\alpha(t)). \end{aligned}$$

Consider the region $V(t)$ with boundary $\alpha(t) \cup \pi(t)$ plus sides to enclose it. Since $\pi(t)$ is the tangent plane to $\alpha(t)$ at $q(t) = q_i + tv$, the height of the sides needed to enclose this region is bounded by $C_2\gamma_1^2$, where since S is compact, $C_2 > 0$ is a constant depending only on S (C_2 is basically the maximum curvature at a smooth point on S). On the other hand, the length of the boundary of $\alpha(t)$ is bounded by $C_3\gamma_1$ for some constant $C_3 > 0$. Therefore since f_n is divergence-free, for some constant $K > 0$,

$$|T(f_n, \alpha(t))(\alpha(t)) - T(f_n, \pi(t))(\pi(t))| \leq C_2C_3\gamma_1^3 + K/n. \quad (29)$$

The analogous result holds for f (without the K/n term), controlling the fourth term above. By Proposition 5.4.4,

$$\lim_{m \rightarrow \infty} |T(g_m, \pi(t))(\pi(t)) - T(f, \pi(t))(\pi(t))| = 0.$$

As in Proposition 5.5.2, if μ_j^n denotes the component measures of f_n , ν_j^m denotes the component measures of g_m , and μ_j denotes the component measures of f , and letting $U = \cup_{t=0}^{\theta_{\gamma_2}} \pi(t)$, we have

$$\left| \int_0^{\theta_{\gamma_2}} T(f_n, \pi(t))(\pi(t)) - T(g_m, \pi(t))(\pi(t)) dt \right|$$

$$\leq \sum_{j=1}^3 |\xi_j(q)| |\mu_j^n(U) - \nu_j^m(U)| \leq \sum_{j=1}^3 |\xi_j(q)| (|\mu_j^n(U) - \mu_j(U)| + |\mu_j(U) - \nu_j^m(U)|).$$

Taking the limit as $m \rightarrow 0$ gives

$$\limsup_{m \rightarrow \infty} \left| \int_0^{\theta_{\gamma_2}} T(f_n, \pi(t))(\pi(t)) - T(g_m, \pi(t))(\pi(t)) dt \right| \leq \sum_{j=1}^3 |\xi_j(q)| |\mu_j^n(U) - \mu_j(U)|.$$

Therefore

$$\begin{aligned} & \left| \int_0^{\theta_{\gamma_2}} T(f_n, \alpha(t))(\alpha(t)) - T(f, \alpha(t))(\alpha(t)) dt \right| \\ & \leq 2C_2 C_3 \theta_{\gamma_2} \gamma_1^3 + \sum_{j=1}^3 |\xi_j(q)| |\mu_j^n(U) - \mu_j(U)| + K/n. \end{aligned}$$

Plugging back in to Equation (26), we get

$$\begin{aligned} \mathbb{W}_1^{1,1}(T(f_n, S), T(f, S)) & \leq 10C\gamma_1 + 2C\gamma_1^{-2}\gamma_2 + C\gamma_1^{-2}Kn^{-1} + 2CC_2C_3\gamma_1 + C\gamma_1^{-2}\theta_{\gamma_2}^{-1}Kn^{-1} + \\ & C\gamma_1^{-2}\theta_{\gamma_2}^{-1} \sum_{j=1}^3 |\xi_j(q)| |\mu_j^n(U) - \mu_j(U)|. \end{aligned}$$

Take $\gamma_2 = \gamma_1^3$. Then taking γ_1 small makes terms 1, 2, and 4 small. Taking n large makes terms 3 and 5 small and by Lemma 5.2.8 also makes term 6 small. \square

Next we remove the condition that S is contained in the interior of R .

Proposition 5.5.6. *For any $S \in \mathbb{S}(R)$ and any sequence of tiling flows $f_n \in TF_n(R)$ such that $d_W(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$ for some $f \in AF(R)$,*

$$\lim_{n \rightarrow \infty} \mathbb{W}_1^{1,1}(T(f_n, S), T(f, S)) = 0.$$

Proof. We can cover S with finitely many surfaces S_1, \dots, S_M which satisfy the conditions of Proposition 5.4.2 for vectors v_1, \dots, v_k and a threshold $\theta > 0$. We can do this so that $d = \max_{1 \leq j \leq M} \text{diam}(S_j)$ and there is a constant C independent of d such that $M = Cd^{-2}$. Fix $\epsilon > 0$. There exists $\delta > 0$, $N > 0$ such that for all $j = 1, \dots, M$, and all $0 \leq t \leq \delta$, by Proposition 5.4.4,

$$\mathbb{W}_1^{1,1}(T(f, S_j), T(f, S_j(t))) < \epsilon$$

and by Proposition 5.5.4, for $n \geq N$ and $0 \leq t \leq \delta$,

$$\mathbb{W}_1^{1,1}(T(f_n, S_j), T(f_n, S_j(t))) < \epsilon.$$

On the other hand, by Proposition 5.5.5, for all $j = 1, \dots, M$ and all $t > 0$, for n large enough

$$\mathbb{W}_1^{1,1}(T(f_n, S_j(t)), T(f, S_j(t))) < \epsilon.$$

Hence by the triangle inequality, for all $j = 1, \dots, M$

$$\mathbb{W}_1^{1,1}(T(f, S_j), T(f_n, S_j)) < 3\epsilon.$$

By the two-dimensional version of Lemma 5.2.7, for n large enough,

$$\mathbb{W}_1^{1,1}(T(f, S), T(f_n, S)) \leq M(10d^3 + 3\epsilon) \leq 10Cd + 3d^{-2}\epsilon.$$

Taking $d = \epsilon^{1/3}$ would complete the proof. \square

Finally we can prove the main theorem about boundary values that we will refer to later in paper.

Theorem 5.5.7. *For any $S \in \mathbb{S}(R)$, the trace operator*

$$T(\cdot, S) : (AF(R) \cup TF(R), d_W) \rightarrow (\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$$

is uniformly continuous. In particular this holds for $S = \partial R$.

Proof. By Theorem 5.3.4, $(AF(R), d_W)$ is compact. Since $TF_n(R)$ is finite for each n , Theorem 5.3.1 implies that the d_W limit points of $TF(R)$ are contained in $AF(R)$. Therefore $(AF(R) \cup TF(R), d_W)$ is compact.

On the other hand, for any $S \in \mathbb{S}(R)$, Proposition 5.4.6 and Proposition 5.5.6 combine to show that $T(\cdot, S)$ is a continuous map from $(AF(R) \cup TF(R), d_W)$ to $(\mathcal{M}^s(R), \mathbb{W}_1^{1,1})$. Therefore by compactness $T(\cdot, S)$ is uniformly continuous. \square

6 Patching

The main goal of this section is to prove a *patching theorem* (Theorem 6.3.5) which will be an essential tool throughout this paper. We show that if the flows associated to tilings τ_1, τ_2 of \mathbb{Z}^3 are *nearly-constant* (Definition 6.3.3 below) with value $s \in \text{Int}(\mathcal{O})$ (which loosely speaking means the flows associated with τ_1, τ_2 both approximate the constant flow equal to s), then we can remove a bounded piece from τ_1 , and patch it to τ_2 by tiling a thin (cubic) annulus.




Figure 21: Two dimensional schematic for patching.

Equivalently, we want to show that this annular region can be tiled by dimers exactly so that it agrees with τ_1 on one boundary and with τ_2 on the other boundary. To do this, we need a condition to show that a region is tileable.

A general condition for tileability, which works in any dimension, is given by the classical *Hall's matching theorem* ([Hal35], stated here as Theorem 6.1.2), which says a region R is tileable if and only if every $U \subset R$ has certain properties; a U that does not have such properties is a “counterexample” to the condition that implies tileability, as we explain in Section 6.1. In two dimensions, Hall's matching theorem implies that a simple condition on height function differences along the boundary of the region is equivalent to tileability [Fou96]. In three dimensions, we show in Section 6.2 that it is sufficient to show that the region R has no counterexample set U whose boundary is a certain type of *minimal surface*, built out

of squares from the \mathbb{Z}^3 lattice (Corollary 6.2.6). We call surfaces built out of lattice squares *discrete surfaces*.

In Section 6.3 we give the statement of the patching theorem, and explain the main ideas of the proof accompanied by a series of two dimensional figures.

The main new difficulty in higher dimensions is that the counterexample sets U can have more complicated geometry. In two dimensions, the boundary of the counterexample region is a union of curves. In three dimensions it is a union of surfaces. However the fact that these surfaces can be assumed to be in some sense *minimal* gives us some control their geometry. In Section 6.4, we prove some straightforward adaptations of the isoperimetric inequality for discrete surfaces. In Section 6.5, we apply these to get useful bounds on the area growth for minimal discrete surface (Proposition 6.5.1), and show that they “spread out” (Lemma 6.5.3).

In Section 6.6, we prove an ergodic theorem for the flow of a tiling through a coordinate plane (Theorem 6.6.3), and note that tilings sampled from ergodic measures satisfy the conditions of the patching theorem with probability going to 1 as $n \rightarrow \infty$ (Corollary 6.6.5). We show that ergodic measures of any mean current $s \in \mathcal{O}$ exist (Lemma 6.6.1), and prove some bounds for their expected flow through discrete surfaces (Lemma 6.6.7). One of the ideas in the proof of the patching theorem is to use a tiling sampled from an ergodic measure as a “measuring stick” that we compare with the tilings we want to patch.

Equipped with the lemmas from the previous sections, in Section 6.7 we give the proof of the patching theorem (Theorem 6.3.5). Finally in Section 6.8 we give some immediate corollaries of patching for ergodic Gibbs measures (EGMs).

We use the same tools and ideas developed in this section again in Section 8.6 to prove a *generalized patching theorem* (Theorem 8.6.2) where the flow the tilings approximate is not required to be constant, and where the annular region is allowed to have a more general shape. The main results of this paper are two versions of a large deviation principle: one with *soft* boundary conditions (Theorem 8.1.6) and one with *hard* boundary conditions (Theorem 8.2.4). The regular patching theorem proved here (Theorem 6.3.5) is sufficient to prove the LDP with soft boundary conditions, but the generalized version (Theorem 8.6.2) is needed in the final steps to prove the version with hard boundary conditions.

6.1 Hall’s matching theorem and non-tileability

When can a finite region $R \subset \mathbb{Z}^3$ be exactly tiled by dimers, i.e. without any tiles crossing the boundary, and with all cubes covered? This is equivalent to asking: when does a finite subgraph $G \subset \mathbb{Z}^3$ have a perfect matching? A straightforward observation is that for any bipartite graph G with bipartition (A, B) , a necessary condition for G to have a perfect matching is that it is *balanced*, i.e. that $|A| = |B|$. The balanced condition is not sufficient though, see Figure 22. Nonetheless, it turns out there is a very general necessary and sufficient condition which characterizes whether or not a finite bipartite graph G has a perfect matching. In this section we explain these results from the graph point of view where dimers are edges and dimer tilings are perfect matchings. There are two perspectives, both of which can be useful:

- the min-cut, max-flow principle
- Hall’s matching theorem




Figure 22: A region that is balanced (i.e. the number of black squares is equal to the number of white squares) but has no dimer tilings.

We first describe the classical min-cut, max-flow principle. Let $G = (A, B)$ be the bipartition of the graph (A is “even” and B is “odd”). If G has a perfect matching τ , then there is a flow v_τ (the “pretiling flow”) which sends a unit of current from each even vertex $a \in A$ to the odd vertex $b \in B$ it is paired to. Note that v_τ has a source of $+1$ at each $a \in A$ and a sink of $+1$ (or source of -1) at each $b \in B$. The existence of a perfect matching is equivalent to the existence of a flow v_τ with the desired source/sink values and a flow of 0 or 1 on each even-to-odd edge.

A *cut* is a collection of edges in G which, if deleted, separates G into two pieces G_1 and G_2 . Let F_1 be the net total flow sourced in G_1 (i.e. the number of even vertices in G_1 minus the number of odd vertices), let F_2 be the net total flow sourced in G_2 , and let c be the number of cut edges. If G has a perfect matching, we must have $F_1 + F_2 = 0$.

The value F_1 measures the amount of flow that would have to travel across the cut if G has a perfect matching, so we must have $F_1 \leq c$. In other words, if G has a perfect matching, then it must be the case that for *every* cut, the excess flow on either side must be less than the size of the cut. It turns out that this is a sufficient condition too, so if G does *not* have a perfect matching then there is a cut of c edges partitioning G into G_1 and G_2 such that the excess flow F_1 that needs to cross the cut is more than c . In summary:

Theorem 6.1.1 (Min-cut, max-flow principle [FF56]). *A finite bipartite graph G has a perfect matching (a.k.a. dimer tiling) if and only if there is no cut consisting of c edges partitioning G into two sets G_1 and G_2 such that $F_1 > c$.*

In Hall’s matching theorem, we shift our perspective from the cut to the sets in the partition. Instead of looking at sets that are a mixture of even and odd vertices, we consider a set C of only even (resp. only odd) vertices, plus their neighbors

$$N(C) = \{b \in B : a \in C, (a, b) \in E\}.$$

The set $C \cup N(C)$ is analogous to either G_1 or G_2 (without loss of generality G_1) plus some of the endpoints of the edges in the cut. The excess flow is now the number of even vertices (i.e. $|C|$) minus the number of odd vertices (i.e. $|N(C)|$). Hall’s matching theorem is an analog of Theorem 6.1.1 formulated in these terms:

Theorem 6.1.2 (Hall’s matching theorem [Hal35]). *Suppose that $G = (V, E)$ is a finite bipartite graph with bipartition $G = A \cup B$. The graph G admits a perfect matching consisting of $|A|$ edges if and only if for all $C \subset A$,*

$$N(C) = \{b \in B : a \in C, (a, b) \in E\}$$

satisfies $|N(C)| \geq |C|$.

An analogous result holds with A and B switched. If G is balanced (i.e. $|A| = |B|$), then the existence of a perfect matching with $|A|$ edges is equivalent to the existence of a perfect matching of the whole graph G . If G is not balanced (i.e. $|A| \neq |B|$) then G does not have a perfect matching of the whole graph.

Note that if $C \subset A$ satisfies $|N(C)| < |C|$, then the set $U := C \cup N(C)$ has more even than odd vertices, despite having only odd vertices on its boundary within G . Therefore when G is balanced, Theorem 6.1.2 is equivalent to the following:

Theorem 6.1.3. *Suppose that $G = (V, E)$ is a finite bipartite graph with bipartition $G = A \cup B$ with $|A| = |B|$. Then G fails to have a perfect matching if and only if there exists a connected set $U \subset V$ such that $|U \cap A| > |U \cap B|$ but all boundary elements of U (i.e., elements of U that are adjacent to some point in $V \setminus U$) belong to B .*

We call the U from Theorem 6.1.3 a *counterexample to (the condition equivalent to) tileability* or just a *counterexample*. In our context, A and B will always be sets of even and odd vertices in \mathbb{Z}^3 , so for us a *counterexample* to tileability for $R \subset \mathbb{Z}^3$ is any set U that has more even than odd vertices, despite having only odd vertices on its *interior boundary*, which we define to be the set of $x \in U$ that are incident to some $y \in R \setminus U$.

To show that a graph R has a dimer tiling (a.k.a. a perfect matching), we check that it is balanced, and if it is, we have to show that there are no counterexamples. We call the excess flow of a counterexample its *imbalance*, given by

$$\text{imbalance}(U) = \text{even}(U) - \text{odd}(U). \quad (30)$$

Note that if $U \subset R \subset \mathbb{Z}^3$ has only odd vertices on its interior boundary (within R) then $\text{imbalance}(U) > 0$ if and only if U is a counterexample for R .

6.2 Discrete surfaces and minimal counterexamples

As mentioned earlier, it is often intuitively useful to think about perfect matchings as tilings by $2 \times 1 \times 1$ blocks. In this picture, a counterexample set U is a collection of unit cubes, each centered at a point in \mathbb{Z}^3 , and the edges out of it, i.e. its boundary ∂U , is a collection of squares in the translated lattice $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) + \mathbb{Z}^3$. In other words the boundary region is a surface built out of squares from the lattice.

Definition 6.2.1. A *discrete surface* in \mathbb{Z}^3 is a collection of squares from the $(\frac{1}{2}, \frac{1}{2}, \frac{1}{2}) + \mathbb{Z}^3$ lattice.

A discrete surface in \mathbb{Z}^3 is *orientable* if there is a well-defined outward pointing normal vector to the surface. An orientable discrete surface S with a choice of outward pointing normal vector is called *oriented*. For a square $s \subset S$, we call the side of s that the outward normal vector points toward the *outside*. If the outward pointing normal vector to a square in an oriented surface is from even to odd, we color the outside of the square white. Otherwise we color it black.

Definition 6.2.2. An oriented discrete surface S in \mathbb{Z}^3 is *monochromatic* if all the outsides of all the squares in S are black (resp. all are white).

We can rewrite Equation (30) for the imbalance of a counterexample U in terms of the black and white surface area of ∂U . Let (A, B) be the bipartition of \mathbb{Z}^3 into even and odd vertices respectively.

Proposition 6.2.3. *Suppose that R is balanced but not tileable, and that $U \subset R \subset \mathbb{Z}^3$ is a counterexample to tileability. Then*

$$0 < \text{imbalance}(U) = \text{even}(U) - \text{odd}(U) = \frac{1}{6} \left(\text{white}(\partial U) - \text{black}(\partial U) \right). \quad (31)$$

Proof. Define a flow r on \mathbb{Z}^3 such that $r(e) = \frac{1}{6}$ for every dual edge (a.k.a. face) e oriented from even to odd. Then

$$\text{div } r(v) = \begin{cases} -1 & v \text{ is a odd cube} \\ +1 & v \text{ is a even cube} \end{cases}$$

By the divergence theorem, with all edges $e \in \partial U$ oriented out of U ,

$$\begin{aligned} \text{imbalance}(U) &= \text{even}(U) - \text{odd}(U) = \sum_{v \in U} \text{div } r(v) = \sum_{e \in \partial U} r(e) \\ &= \frac{1}{6} \left(\text{white}(\partial U) - \text{black}(\partial U) \right). \end{aligned}$$

□

By Proposition 6.2.3, if U is a counterexample then it must have more white surface area than black surface area. By the definitions in Section 6.1, U must have only odd cubes along its *interior boundary*, i.e. cubes $x \in U$ which are adjacent to $y \in R \setminus U$. However U also has an *exterior boundary* consisting of cubes which are adjacent to $y \in \mathbb{Z}^3 \setminus R$. Exterior boundary cubes can be even or odd.

Correspondingly, the boundary ∂U can be split into two pieces: the *exterior boundary surface* $T = \partial R \cap \partial U$ and the *interior boundary surface* $S = \partial U \setminus T$. The interior boundary surface S must be built out of only black squares, while T can be built out of a mixture of white and black squares. Given this, only squares in T contribute positively to the imbalance of U . Intuitively to increase the imbalance of U , one should *minimize* the area of S .

A surface P embedded in \mathbb{R}^3 is said to *locally minimize area* if given any point $p \in P$, there is a neighborhood $V \subset \mathbb{R}^3$ containing p such that $P \cap V$ has the minimal area of any surface with boundary $\partial(P \cap V)$. Surfaces that locally minimize area are called *minimal surfaces*. We will be interested in certain discrete analogs of minimal surfaces.

Definition 6.2.4. A *minimal discrete (monochromatic) surface with boundary X* is a surface S that minimizes area subject to the constraint that it is discrete, (monochromatic), and has $\partial S = X$. In particular, there is no way to “tighten the surface locally” by changing a few faces in a way that maintains the overall boundary conditions and reduces the overall area.

Proposition 6.2.5. *Let $R \subset \mathbb{Z}^3$ be a finite balanced region which is not tileable, and suppose that U is a counterexample to tileability in R . Let $T = \partial U \cap \partial R$ and let $S = \partial U \setminus T$ be the interior boundary surface. Then there exists another counterexample U' in R such that $\partial U' \cap \partial R = T$, and $S' := \partial U' \setminus T$ is a minimal monochromatic black discrete surface.*

Proof. The new set U' is defined so that $\partial U' = T \cup S'$, where S' is a minimal monochromatic discrete surface. Since S is all black, by Proposition 6.2.3,

$$6 \cdot \text{imbalance}(U) = \text{white}(\partial U) - \text{black}(\partial U) = \text{white}(T) - \text{black}(T) - \text{area}(S) \quad (32)$$

replacing S by S' only makes the imbalance larger, so U' is still a counterexample. \square

We call counterexamples where the internal boundary surface is a minimal surface *minimal counterexamples*. We immediately get the following corollary.

Corollary 6.2.6. *A finite balanced region $R \subset \mathbb{Z}^3$ is tileable if and only if it has no minimal counterexamples.*

6.3 Statement of patching theorem and outline of proof

We now give the statement of the patching theorem mentioned at the beginning of the section. We also provide illustrations of the analogous constructions in 2D (because they are easier to draw) in order to explain the 3D concepts. Let $B_n = [-n, n]^3$, and for any $\delta > 0$ define

$$A_n = B_n \setminus B_{(1-\delta)n}.$$

Given two tilings $\tau_1, \tau_2 \in \Omega$, we look at the region between $\tau_1|_{\mathbb{Z}^3 \setminus B_n}$ and $\tau_2|_{B_{(1-\delta)n}}$. This will be the annulus A_n , with some cubes removed along its boundary (see Figure 23). We call this A_n with boundary conditions τ_1 and τ_2 .




Figure 23: Here is an example of a region of the form “ A_n with boundary conditions τ_1 and τ_2 ” in two dimensions, which we call A . The region shown is a subset of the “square annulus” bounded between the boundary of a 12×12 box and the boundary of a 32×32 box. It is obtained from the square annulus by removing some of the squares along the outer boundary and some of the squares along the inner boundary. Given any dimer tiling τ of a region containing the 32×32 box, the union of the dimers of τ that are *strictly contained* in the square annulus would also be a region of this form.

The main question is: given a tiling τ_1 restricted to $\mathbb{Z}^3 \setminus B_n$ and a tiling τ_2 restricted to $B_{(1-\delta)n}$, under what conditions can we patch them together, i.e. find a tiling of A_n with inner boundary condition τ_2 and outer boundary condition τ_1 ? We are interested in showing that this is possible when n is large enough, when τ_1, τ_2 satisfy a consistency condition that they are *nearly constant* for the same $s \in \text{Int}(\mathcal{O})$.

To specify the nearly constant condition, we give a few definitions.

Definition 6.3.1. An ϵ patch α on ∂B_n is an $\epsilon n \times \epsilon n$ square contained in a face of ∂B_n .

We can then measure the *flux* of a discrete flow through a patch.

Definition 6.3.2. Let S be an oriented discrete surface with outward normal vector ξ . We define the *flux* of a discrete vector field v through S by

$$\text{flux}(v, S) = \sum_{e \in E(\mathbb{Z}^3), e \cap S \neq \emptyset} \text{sign}\langle \xi(e \cap S), e \rangle v(e).$$

Here $E(\mathbb{Z}^3)$ denotes the edges of \mathbb{Z}^3 oriented from even to odd.

We now use the definition of *nearly-constant* in terms of the flux of the pre-tiling flow v_τ through patches.

Definition 6.3.3. Fix $s \in \mathcal{O}$, let $B_n = [-n, n]^3$. A tiling $\tau \in \Omega$ is ϵ -*nearly-constant with value* s if there exists $M = M(\epsilon)$ such that for all $n > M$ and all ϵ patches α on ∂B_n ,

$$\text{flux}(v_\tau, \alpha) = \frac{1}{2} \langle \xi_\alpha, s \rangle \text{area}(\alpha) + o(\text{area}(\alpha)) = \frac{1}{2} \langle \xi_\alpha, s \rangle \epsilon^2 n^2 + o(\epsilon^2 n^2),$$

where ξ_α is the outward pointing unit normal vector to α (where outward means away from B_n).

Remark 6.3.4. Any patch α is contained in a flat coordinate plane, so its area is simply $\epsilon^2 n^2$. The $\frac{1}{2}$ comes from the fact that the mean current is actually measuring the average flow per *even* vertex.

Tilings satisfying the ϵ -nearly-constant condition with value s mimic the behavior of tilings sampled from ergodic measures of mean current s . (This is made precise in Corollary 6.6.5 after Theorem 6.6.3.)

With the conditions defined, we can now state the patching theorem.

Theorem 6.3.5 (patching theorem). *Fix $\delta > 0$ and a mean current $s \in \text{Int}(\mathcal{O})$. Let $B_n = [-n, n]^3$ be the cube of radius n , and let $A_n = B_n \setminus B_{(1-\delta)n}$ be the cube annulus of width δn . For $\epsilon > 0$ small enough, if $\tau_1, \tau_2 \in \Omega$ are ϵ -nearly-constant with value s , then for n large enough A_n can be tiled with outer boundary condition τ_1 and inner boundary condition τ_2 .*

The main tool in the proof is Hall's matching theorem (Theorem 6.1.2). In this section we explain the main steps of the proof, guided by a series of two dimensional figures, and comment on differences between the two and three dimensional versions of this story. After this, in the remaining subsections we prove a series of lemmas (needed to control the more complicated geometric situations that can occur in three dimensions) before giving a the formal proof of Theorem 6.3.5 in Section 6.7.

Steps of proof

Given two tilings τ_1, τ_2 of \mathbb{Z}^2 , we want to know whether they can be patched together. In other words, we want to know whether a region A , which is a square annulus with some squares removed along the outer boundary if they are connected by τ_1 to the outside of the annulus and some squares removed along the inner boundary if they are connected by τ_2 to the outside of the annulus, is tileable. See Figure 23 for an example of a region A of this form. If A is not tileable, then there exists a *counterexample set* U as in Figure 24.




Figure 24: A potential counterexample set $U \subset A$.

The red set $U \subset A$ in Figure 24 has the property that every square on its *inner boundary* (i.e., every square of U that is incident to a square in $A \setminus U$) is black.

By Hall’s matching theorem (Theorem 6.1.3), there exists a dimer tiling of A if and only if every U of this form is *not* a counterexample. In other words, every set U of this form has $\text{imbalance}(U) = \text{white}(U) - \text{black}(U) \leq 0$. We remark that the colors white and black used here stand for even and odd vertices and not for the colors we give to surfaces in 3 dimensions in the previous section. We do this because it becomes easier to illustrate the main ideas using the figures.

Given this, our strategy to show that A is tileable for n large enough proceeds by contradiction. We suppose that for all n there is a set U of the form above that is a counterexample, but then show that for n large enough, it must have $\text{white}(U) \leq \text{black}(U)$ and hence not be a counterexample. To do this, we cut up U into various smaller pieces, and bound the white minus black in each piece.

First we divide the annular A into regions as depicted in Figure 25. We call these the “middle region” (blue), the “thin region” (orange) and the “corner region” (green). The middle region is a centered square annulus whose size will have to be appropriately tuned. The thin region is the union of the “columns” obtained as straight-line paths of squares that go from the middle square annulus boundary to the boundary of A . The corner region is the part leftover.

We then define U' to be U intersected with the middle layer, depicted in Figure 26.




Figure 25: The middle (blue), thin (orange), and corner (green) regions of A .




Figure 26: We can define U' to be the intersection of U with the middle region from Figure 25.

Given a tiling τ of a region containing A , we can define U_τ to be the region covered by tiles from τ intersecting U' . The set U'' shown in Figure 27 is the subset of U_τ that consists of the union of U' together with all of the squares covered by tiles from τ that are contained in U but in the complement of the middle region. Note that $U_\tau \setminus U''$ consists of only white squares and that U_τ is by construction evenly balanced between black and white squares. If we can show that $|U_\tau \setminus U''|$ is large then we know that U'' has many more black squares than white.

Indeed, we show that we can choose this *test tiling* τ so that $|U_\tau \setminus U''|$ is large and hence U'' has many more black than white squares. Since $U'' \subset U$, it remains to show that there cannot be enough white squares in $U \setminus U''$ for U to be a counterexample.




Figure 27: The region U'' (blue) and the tiles from τ that intersect U' (green).

In order to prove that U itself has more black than white squares, we will divide the rest of U into multiple pieces to treat separately, depicted in Figure 28. Here U'' is as given in Figure 27. The “shadow region” is $U \setminus U''$ restricted to the “thin region” from Figure 25. It consists of the union of the columns that can be extended all the way from ∂A to $\partial U''$. The corner region here is the intersection of U with the corner region from Figure 25. The “leftover pieces” are the parts of U that do not belong to one of the other three regions. Roughly speaking, we aim to show that U has more black than white by showing that (i) U'' has a lot more black than white, (ii) the shadow region can only have a little more white than black (because of the nearly-constant condition), (iii) the corner region can only have a little excess white (since it has small volume), and (iv) the leftover pieces have at least as much black as white.




Figure 28: Depiction of all the regions that we divide U into: (i) U'' , (ii) shadow region, (iii) corner region, and (iv) leftover pieces.

Ultimately to achieve (i), (ii), (iii), and (iv), we need to understand more about the possible shapes of counterexample sets U . The main difference between two dimensions and three dimensions is that the geometry of U can be more complicated.

In two dimensions, since A is balanced, if U is a counterexample then so is U^c (with even/odd switched). As a consequence, without loss of generality U and U^c are both connected, and (since U and U^c must both cross the annulus) the interior boundary curve S between them consists of two minimal monochromatic paths S_1, S_2 from the inner boundary to the outer boundary. There are then four points a_1, b_1 and a_2, b_2 which are the endpoints of S_1, S_2 on the inner and outer boundaries respectively. From here, one can argue as in [Fou96] to say that A is tileable as long as the boundary height differences (given by the boundary condition tilings τ_1, τ_2) at a_1 and b_1 and at a_2 and b_2 are in some sense compatible.

In three dimensions, since A is balanced, it again follows that both U and U^c are counterexamples, and that they are without loss of generality connected and have interior boundary surface S between them which is a minimal monochromatic discrete surface. However the geometry (and topology) of 2D surfaces and connected sets S in 3D can be much more complicated than that of 1D curves in 2D. For example, the set U in 3D need not be simply connected, and instead of $S \cap \partial A$ consisting of four points, it consists of a collection of curves.

To control the more complex *geometric possibilities* for counterexamples in three dimensions, we prove that the interior boundary surface S has area of order n^2 (Lemma 6.5.2), and that there is a choice for the inward blue layer where S restricted to the layer (which is some union of curves) has good properties (Lemma 6.5.3). This is the content of Section 6.4 and Section 6.5.

The other key tool is the notion of a *test tiling*. We show that we can use a tiling τ sampled from an ergodic measure of mean current $s \in \text{Int}(\mathcal{O})$ to define U'' so that U'' will have order n^2 more black cubes than white cubes by showing that τ on expectation has an order n^2 number of tiles crossing S (Lemma 6.6.7). This quantity is somewhat analogous to the height

difference in two dimensions. Further, a tiling sampled from ergodic measure of mean current $s \in \text{Int}(\mathcal{O})$ is nearly constant with high probability (Corollary 6.6.5). This is the content of Section 6.6.

6.4 Discrete isoperimetric inequalities

The classical isoperimetric inequality says that the minimal area of a region D bounded by a smooth closed curve γ in \mathbb{R}^2 of length l is $\frac{1}{4\pi}(l^2)$. The equality case is achieved when γ is a circle and D is a disk. A lesser known fact is that this bound also holds for curves in \mathbb{R}^3 [Alm86]. In this section we prove a discrete version of the isoperimetric inequality.

Proposition 6.4.1. *Given any simple closed curve γ in \mathbb{Z}^3 , there is a surface S with $\partial S = \gamma$ and*

$$\text{area}(S) \leq \frac{1}{8}(\text{length}(\gamma))^2.$$

Remark 6.4.2. The constant $\frac{1}{8}$ that we get from the proof is not optimal but our argument is simple and the result is sufficient for our purposes. Also note that the boundary curve γ can be replaced by a multicurve $\Gamma = \gamma_1 \cup \dots \cup \gamma_k$ and the same result holds.

Proof. We proceed by induction. Suppose that for all simple closed curves γ in \mathbb{Z}^3 with $\text{length}(\gamma) = 2m \leq 2n$, there exists a surface S with $\partial S = \gamma$ and

$$\text{area}(S) \leq \frac{(m-1)m}{2}.$$

This is sufficient because it implies that

$$\text{area}(S) \leq \frac{1}{8}(\text{length}(\gamma))^2.$$

Now suppose that β is a simple closed curve of length $2n+2$, and equip β with an orientation. Choose two parallel edges $e_1 = (a_1, b_1), e_2 = (a_2, b_2)$ in β with opposite orientations. Removing e_1, e_2 and identifying a_i with b_i for $i = 1, 2$ results in a new curve β' of length $2n$. The identification means that β' is a union simple curves $\gamma_1, \dots, \gamma_k$ of length $\leq 2n$ and double edges (note that each double edge contributes 2 to the length of β'). By the inductive hypothesis, there exist surfaces S_i with boundary γ_i satisfying the bound for each $i = 1, \dots, k$. To get a surface with boundary β , we find a path in \mathbb{Z}^3 from a_1 to a_2 along $\cup_{i=1}^k S_i \cup D$, where D is the double edges in β' . Since $\text{length}(\beta') \leq 2n$, we can find a path of length $\leq n$. We add back the edges e_1, e_2 , splitting the path into two parallel paths. We then add squares along the path from e_1 to e_2 to construct a surface S with boundary β satisfying

$$\text{area}(S) \leq \text{area}(\cup_{i=0}^k S_i) + n \leq \frac{n(n+1)}{2}.$$

□

This can easily be extended to the monochromatic case; the only thing that changes is the constant.

Corollary 6.4.3. *Given any collection of simple closed curves $\gamma_1, \dots, \gamma_k$ in \mathbb{Z}^3 , there is a monochromatic surface S with $\Gamma = \gamma_1 \cup \dots \cup \gamma_k$ as its boundary and*

$$\text{area}(S) \leq \frac{5}{8}\text{length}(\Gamma)^2$$

Proof. Using Proposition 6.4.1, we can find a surface T with boundary $\gamma_1 \cup \dots \cup \gamma_k$ satisfying $\text{area}(T) \leq \frac{1}{8} \text{length}(\Gamma)^2$. Replacing every white square in T by at most 5 black squares, we get a new surface S which is monochromatic (it is all black) satisfying the same bound with the constant $\frac{5}{8}$. \square

6.5 Area growth of minimal monochromatic discrete surfaces

We now use the discrete isoperimetric inequalities from the previous section to show that minimal monochromatic discrete surfaces have quadratic area growth (Proposition 6.5.1). We then apply this to the cube annulus $A_n = B_n \setminus B_{(1-\delta)n}$ to get two results (Lemma 6.5.2 and Lemma 6.5.3) which will serve as lemmas for the patching theorem (Theorem 6.3.5).

Proposition 6.5.1. *Let S be a minimal monochromatic discrete surface. Let $p \in \mathbb{Z}^3 + (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ be a vertex on S and let $B_n(p) = p + [-n, n]^3$ be such that S is not contained in $B_n(p)$. Then there is a universal constant $\kappa > 0$ (i.e., independent of S and n) such that*

$$\text{area}(S) \geq \kappa n^2.$$

Proof. To do this, we show that $\text{area}(S \cap B_n(p))$ grows quadratically in n . Let $m = \lfloor \frac{n}{2} \rfloor$. For $k \leq m$, define annular regions

$$A_k = B_{m+k}(p) \setminus B_{m-k}(p)$$

Let $S_k = S \cap A_k$ be S restricted to A_k and let $\Gamma_k = \partial S_k$. By Corollary 6.4.3,

$$\text{area}(S_k) \leq \frac{5}{8} \text{length}(\Gamma_k)^2.$$

Note that $S_k \cap \partial A_k$ might be larger than Γ_k , since there might be squares from S_k contained in ∂A_k .

Any face in $S_{k+1} \setminus S_k$ corresponds to at most 4 edges along Γ_k . Therefore

$$\text{area}(S_{k+1}) - \text{area}(S_k) \geq \frac{1}{4} \text{length}(\Gamma_k) \geq \frac{1}{4} \sqrt{\frac{5}{8}} \sqrt{\text{area}(S_k)}.$$

Therefore the function $f(k) := \text{area}(S_k)$ satisfies the inequality $f(k+1) - f(k) \geq \sqrt{f(k)}$. Extending f linearly to a continuous function that is differentiable from the left, this becomes $f'(k) \geq \sqrt{f(k)}$, which implies that f grows at least quadratically in k . Applying this to S itself we get that $\text{area}(S) \geq \kappa n^2$, where the universal constant κ comes from the isoperimetric inequalities. \square

Lemma 6.5.2. *Let $B_n = [-n, n]^3$ and suppose that $A = B_n \setminus B_{(1-\delta)n}$ for some $\delta \in (0, 1)$. Suppose that S is a minimal monochromatic discrete surface in A which connects the inner and outer boundaries of A . Then there exist constants c_1, c_2 independent of S and n such that*

$$c_1 n^2 \geq \text{area}(S) \geq c_2 n^2$$

where $c_2 \sim \delta^2$.

Proof. For the upper bound, notice that $\partial S \subset \partial A$. From ∂A , we construct a monochromatic surface T by capping every white square on ∂A with a odd cube. There is a surface $S' \subset T$ with the same boundary as S . Since $\text{area}(T) \leq 5 \text{area}(\partial A)$, there exists a constant c_1 such that $\text{area}(T) < c_1 n^2$. Since S is a minimal monochromatic surface,

$$c_1 n^2 \geq \text{area}(S') \geq \text{area}(S).$$

Now we prove the lower bound. Since S connects the inner and outer boundaries, there is a point $p \in S$ where we can apply Proposition 6.5.1 to $B_{(\delta/3)n}(p)$. Hence

$$\text{area}(S) \geq \kappa((\delta/3)n)^2 =: c_2 n^2$$

and c_2 is of order δ^2 . □

The next application of the area growth results is loosely that minimal surfaces “spread out.” If $X \subset \partial B_n$ for some n is a surface, we define the ϵ covering area of X to be the total area in disjoint $\epsilon n \times \epsilon n$ size squares needed to contain X . We denote this by $\text{Cov}_\epsilon(X)$.

Lemma 6.5.3 (Indenting lemma). *Fix $\delta > 0$. Let $A = B_n \setminus B_{(1-\delta)n}$, and suppose that S is a minimal monochromatic discrete surface connecting the inner and outer boundaries of A . Let $\Gamma_l = \partial(S \cap \partial B_l)$. There exist constants $c, c' > 0$ independent of S and n such that for any $\epsilon > 0$, there exists $l \in ((1 - \epsilon^{1/2})n, n)$ such that*

$$\text{Cov}_\epsilon(S \cap \partial B_l) \leq c \epsilon^{1/2} n^2$$

and

$$\text{length}(\Gamma_l) \leq c' \epsilon^{-1/2} n.$$

The analogous statements also hold for $l \in ((1 - \delta)n, (1 - \delta + \epsilon^{1/2})n)$.

Proof. Cut the region between ∂B_n and $\partial B_{(1-\epsilon^{1/2})n}$ into $M = \lfloor \frac{1}{3\epsilon^{1/2}} \rfloor$ layers L_1, \dots, L_M of width $\frac{\epsilon^{1/2}n}{M} \geq 3\epsilon n$. By Lemma 6.5.2 plus the pigeonhole principle, there exists j such that

$$\text{area}(S \cap L_j) \leq \frac{c_1 n^2}{M} \leq 3c_1 \epsilon^{1/2} n^2. \quad (33)$$

We further subdivide L_j into three layers $L_j^{(1)}, L_j^{(2)}, L_j^{(3)}$, where $L_j^{(2)}$ is the middle one. These each have width at least ϵn . The l we find satisfying the conditions will have $\partial B_l \subset L_j^{(2)}$.

Let J be the number of $(\epsilon n)^3$ sized boxes needed to cover $S \cap L_j^{(2)}$. Then there are at least $J/9$ disjoint cubes of size $(3\epsilon n)^3$ in L_j , centered on an $(\epsilon n)^3$ cube in $L_j^{(2)}$, such that the central cube is in J (i.e., S intersects the $(\epsilon n)^3$ central cube of this $(3\epsilon n)^3$ cube). Given this, we can apply Proposition 6.5.1 to a point in each of the $J/9$ boxes with radius ϵn to get the bound

$$\text{area}(S \cap L_j) \geq \kappa \frac{J}{9} (\epsilon n)^2. \quad (34)$$

Combining Equations (33) and (34) and solving for J gives

$$J \leq \frac{27c_1}{\kappa} \epsilon^{-3/2}.$$

Therefore for any l such that $\partial B_l \subset L_j^{(2)}$,

$$\text{Cov}_\epsilon(S \cap \partial B_l) \leq J(\epsilon n)^2 \leq \frac{27c_1}{\kappa} \epsilon^{1/2} n^2.$$

so the first part holds with $c = 27c_1/\kappa$.

For the second part, we note that $L_j^{(2)}$ has width ϵn . Any square in $S \cap L_j^{(2)}$ contributes length at most four to the curves Γ_l for l such that $\partial B_l \subset L_j^{(2)}$, thus $\sum_{l: \partial B_l \subset L_j^{(2)}} \text{length}(\Gamma_l) \leq 4 \text{area}(S \cap L_j^{(2)})$. Therefore by Equation (33) and the pigeonhole principle again, we can find l with $\partial B_l \subset L_j^{(2)}$ such that

$$\text{length}(\Gamma_l) \leq \frac{12c_1 \epsilon^{1/2} n^2}{\epsilon n} = 12c_1 \epsilon^{-1/2} n,$$

so the second part holds with $c' = 12c_1$. \square

6.6 Tilings sampled from ergodic measures

Recall that \mathcal{P}_e denotes the measures on Ω which are ergodic with respect to the action of $\mathbb{Z}_{\text{even}}^3$, and that $\{\eta_i\}_{i=1}^3$ denote the standard unit basic vectors. In this section we prove a few results for tilings sampled from ergodic measure of mean current $s \in \mathcal{O}$. In the proof of the patching theorem, we use a ‘‘test tiling’’ sampled from an ergodic measure which we compare with the two other tilings we want to patch. First we note that there exist ergodic measures of mean current s for all $s \in \mathcal{O}$.

Lemma 6.6.1. *For every $s \in \mathcal{O}$, there exists an ergodic measure on dimer tilings of \mathbb{Z}^3 of mean current s .*

Remark 6.6.2. We use methods called *chain swapping* described in Section 7.4 to construct ergodic measures of any edge, then face, then interior mean current from the ones for $s \in \mathcal{V}$. The only results that we use about chain swapping here (Propositions 7.4.2 and 7.4.4) are essentially computations, and do not rely on any results presented in this section.

Note that by Theorem 4.2.7, there exist ergodic measures for all mean currents $s \in \partial\mathcal{O}$, so we only need to use chain swapping to show existence for $s \in \text{Int}(\mathcal{O})$. However we choose not to rely on this here, since the chain swapping techniques allow us to show existence easily just from existence of ergodic measures at the vertices of $\partial\mathcal{O}$.

Proof of Lemma 6.6.1. Let $\mathcal{V} \subset \partial\mathcal{O}$ denote its vertices. For each $s \in \mathcal{V}$, the atomic measure which samples the corresponding brickwork pattern is an ergodic measure of mean current s .

Given any $s \in \text{Int}(\mathcal{O})$ (resp. s contained in a face of $\partial\mathcal{O}$, resp. an edge of $\partial\mathcal{O}$), there exists $p \in (0, 1)$ such that

$$s = (1 - p)s_1 + ps_2$$

for $s_1, s_2 \in \partial\mathcal{O}$ (resp. contained in the edges of $\partial\mathcal{O}$, resp. contained in $\mathcal{V} \subset \partial\mathcal{O}$). Let μ_1 and μ_2 be ergodic measures of mean current s_1, s_2 respectively. Let μ be a measure on $\Omega \times \Omega$ which is an ergodic coupling of μ_1 and μ_2 , let μ' be obtained from μ by chain swapping with

swap probability p , and let μ'_1 and μ'_2 denote its marginals. By Proposition 7.4.4, the mean current of μ'_1 is

$$s(\mu'_1) = (1 - p)s_1 + ps_2 = s.$$

By Proposition 7.4.2, μ' is an ergodic measure on $\Omega \times \Omega$. Therefore μ'_1 is an ergodic measure of mean current s . \square

Recall that the *pretiling flow* v_τ is defined for e oriented from even to odd by

$$v_\tau(e) = \begin{cases} +1 & e \in \tau \\ 0 & e \notin \tau. \end{cases}$$

Let S be an oriented discrete surface with outward normal vector ξ . Applying Definition 6.3.2 to v_τ , the *flux* of v_τ through S by

$$\text{flux}(v_\tau, S) = \sum_{e \in E(\mathbb{Z}^3), e \cap S \neq \emptyset} \text{sign}\langle \xi(e \cap S), e \rangle v_\tau(e).$$

As in the definition, $E(\mathbb{Z}^3)$ denotes the edges of \mathbb{Z}^3 oriented from even to odd. Flux of v_τ has a simple combinatorial interpretation. It counts the number of tiles in τ which cross S , with sign corresponding to the parity of the tile. If S is monochromatic black, then $\text{flux}(v_\tau, S)$ is minus the number of tiles in τ which cross S .

Theorem 6.6.3. *Let P be a coordinate plane with normal vector η_i for some $i = 1, 2, 3$, and let $P_n = P \cap B_n$ (recall $B_n = [-n, n]^3$). If $\mu \in \mathcal{P}_e$ has mean current $s \in \mathcal{O}$, then*

$$\lim_{n \rightarrow \infty} \frac{1}{|P_n|} \text{flux}(v_\tau, P_n) = \frac{1}{2} \langle s, \eta_i \rangle$$

where this limit converges almost surely and in probability.

Remark 6.6.4. The completely equivalent statement holds with v_τ replaced by f_τ . In fact since P_n is contained in a coordinate plane $\text{flux}(f_\tau, P_n) = \text{flux}(v_\tau, P_n) + o(1)$. The reason for the $\frac{1}{2}$ factor is that the mean current is the average current per *even* vertex. The number of even vertices in P_n is $|P_n|/2$.

Proof. Without loss of generality let P denote the (x, y) coordinate plane, so the normal vector is $\eta_3 = (0, 0, 1)$. Recall that discrete surfaces consist of squares in $\mathbb{Z}^3 + (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, so $P = \{(x, y, 1/2) : x, y \in \mathbb{R}\}$. Let $\mathbb{Z}_{\text{even}}^2 = (\mathbb{Z}^2 \times \{0\}) \cap \mathbb{Z}_{\text{even}}^3$. Since μ is invariant under the $\mathbb{Z}_{\text{even}}^2$ action, we can apply the ergodic theorem for this subaction. Let

$$T = P \cap ([-1/2, 3/2] \times [-1/2, 1/2] \times [0, 1])$$

and $\text{even}(B_n) = \mathbb{Z}_{\text{even}}^3 \cap B_n$. The set T is defined so that it contains the intersection points of two adjacent edges, namely $(0, 0, 0)$ to $(0, 0, 1)$ and $(1, 0, 0)$ to $(1, 0, 1)$. Consider the function $F : \Omega \rightarrow \mathbb{R}$ given by $F(\tau) = \frac{1}{2} \text{flux}(v_\tau, T)$. We have that

$$\left| \frac{1}{|\text{even}(B_n)|} \sum_{\eta \in B_n} F(\tau + \eta) - \frac{1}{|P_n|} \text{flux}(v_\tau, P_n) \right| = o(n^2)$$

and by the ergodic theorem

$$\lim_{n \rightarrow \infty} \frac{1}{|P_n|} \text{flux}(v_\tau, P_n)$$

exists in probability and almost surely. Temporarily we call the limit $\text{flux}^*(v_\tau, P)$. We know that this is $\mathbb{Z}_{\text{even}}^2$ -invariant. By integrating the flux across T we get that

$$\int_{\Omega} \text{flux}^*(v_\tau, P) \, d\mu(\tau) = \int_{\Omega} F(\tau) \, d\mu(\tau) = \frac{1}{2} \langle s, \eta_3 \rangle.$$

If we can now prove that the average $\text{flux}^*(v_\tau, P)$ is in fact not just $\mathbb{Z}_{\text{even}}^2$ -invariant but also $\mathbb{Z}_{\text{even}}^3$ -invariant it will follow that it is constant almost surely and equal to $\frac{1}{2} \langle s, \eta_3 \rangle$. To show this, we use the fact that v_τ is essentially “divergence free”. Indeed the tiling flow $f_\tau = v_\tau - r$ where r is a reference flow and the flux of r across T is zero. It follows that

$$|\text{flux}(f_\tau, P_n) - \text{flux}(v_\tau, P_n)| = o(n)$$

and thereby we have that

$$\lim_{n \rightarrow \infty} \frac{1}{|P_n|} \text{flux}(v_\tau, P_n) = \lim_{n \rightarrow \infty} \frac{1}{|P_n|} \text{flux}(f_\tau, P_n).$$

Since f_τ is divergence free it follows that for all $\gamma \in \mathbb{Z}_{\text{even}}^3 \setminus \mathbb{Z}_{\text{even}}^2$,

$$|\text{flux}(f_\tau, P_n) - \text{flux}(f_{\tau+\gamma}, P_n)|$$

is given by the flux through the sides of parallelepiped formed by P_n and $P_n + \gamma$. Thus for a fixed $\gamma \in \mathbb{Z}_{\text{even}}^3$

$$|\text{flux}(f_\tau, P_n) - \text{flux}(f_{\tau+\gamma}, P_n)| = o(n)$$

and hence the same holds for v_τ in place of f_τ . Thus $\text{flux}^*(v_\tau, P)$ is $\mathbb{Z}_{\text{even}}^3$ -invariant, which completes the proof. \square

As a straightforward corollary of Theorem 6.6.3, we see that tilings sampled from ergodic measures satisfy the ϵ -nearly-constant condition with high probability.

Corollary 6.6.5. *Fix $\epsilon > 1$. If μ is an ergodic measure of mean current s , then a tiling τ sampled from μ is ϵ -nearly-constant on B_n with value s with probability arbitrarily close to 1 for n large enough.*

The final goal of this section is to get an estimate on the expected flux of a pretiling flow across a monochromatic discrete surface (e.g. the boundary of a counterexample). To do this, we use the following combinatorial result.

Lemma 6.6.6. *Suppose that S is a monochromatic discrete surface. Let $X_1, X_2, X_3 \subset S$ be the sets of squares with normal vectors η_1, η_2, η_3 respectively.*

If S is a closed surface, then $|X_1| = |X_2| = |X_3|$. If S has boundary ∂S , then for all pairs $i \neq j$,

$$|X_i| = |X_j| + O(\text{length}(\partial S)).$$

Proof. If S is a closed discrete surface, then every edge of S is contained in either two or four squares from S (four can happen if the edge is an edge of non-manifold points). For an edge contained in two squares from the surface, we say those squares are *neighbors*. For an

edge contained in four squares, we arbitrarily split the four into pairs, and say that the paired squares are neighbors. With this definition, every square $f \in S$ has exactly four neighbors. Since S is monochromatic, if $f \in X_1$, then two of its neighbors are in X_2 and two of its neighbors are in X_3 (and similarly for any permutation of 1, 2, 3).

View the set of all squares as a graph, where each square corresponds to a vertex, and two vertices are connected by an edge if the corresponding squares are neighbors. The number of edges connecting X_1 to X_2 must be equal to $2|X_1|$ (since every $f \in X_1$ has two neighbors in X_2), and analogously must be equal to $2|X_2|$ (since every $f \in X_2$ has two neighbors in X_1). Therefore $|X_1| = |X_2|$. An analogous argument shows that $|X_3| = |X_1|$ and completes the proof in the closed surface case.

If S is not closed, then squares $f \in S$ which contain an edge along ∂S do not have exactly four neighbors. Therefore the result holds up to an error of $\text{length}(\partial S)$. \square

Lemma 6.6.7. *Let S be a monochromatic black surface with boundary ∂S , and let Θ be the collection of odd cubes adjacent to S . For any tiling τ ,*

$$|\text{flux}(v_\tau, S)| \leq |\Theta|.$$

If μ is an ergodic measure of mean current $s \in \text{Int}(\mathcal{O})$, then there is a constant $K_\mu \in (0, 1)$ independent of S such that

$$\mathbb{E}_\mu[|\text{flux}(v_\tau, S)|] \geq K_\mu |\Theta| + O(\text{length}(\partial S)).$$

Remark 6.6.8. If μ has mean current $s \in \partial\mathcal{O}$, then $K_\mu = 0$.

Remark 6.6.9. Since S is monochromatic black, $\text{flux}(v_\tau, S)$ is minus the number of tiles from τ crossing S . This is why we add the absolute value.

Proof. Any tile $e \in \tau$ crossing S contains a cube from Θ . From this it follows immediately that $|\text{flux}(v_\tau, S)| \leq |\Theta|$.

Let p_1, \dots, p_6 be the probabilities under μ for the six types of tiles, ordered so that $s_1 = p_1 - p_2$, $s_2 = p_3 - p_4$ and $s_3 = p_5 - p_6$. Similarly let N_1, \dots, N_6 be the six types of squares on S , where the tile type parallel to the outward pointing normal vector at a square $f \in N_i$ has probability p_i . The random variable $|\text{flux}(v_\tau, S)|$ can also be written as $\sum_{f \in S} \mathbb{1}_f(\tau)$, where $\mathbb{1}_f(\tau)$ is the indicator variable which is 1 if there is a tile in τ crossing f and 0 otherwise. From this, we see that

$$\mathbb{E}_\mu[|\text{flux}(v_\tau, S)|] = \sum_{i=1}^6 p_i N_i.$$

We minimize the right hand side to get a positive lower bound (clearly 0 is a lower bound). Let $\text{area}(S) = N = \sum_{i=1}^6 N_i$. By Lemma 6.6.6, $N_1 + N_2$, $N_3 + N_4$ and $N_5 + N_6$ are equal to $N/3$ up to an error of $O(\text{length}(\partial S))$. Thus

$$\begin{aligned} \sum_{i=1}^6 p_i N_i &\geq \max\{p_1 N_1 + p_2 N_2, p_3 N_3 + p_4 N_4, p_5 N_5 + p_6 N_6\} \\ &\geq \max\{\min\{p_1, p_2\}, \min\{p_3, p_4\}, \min\{p_5, p_6\}\}(N/3 + O(\text{length}(\partial S))). \end{aligned}$$

Since $s \in \text{Int}(\mathcal{O})$, at least four of $\{p_i\}_{i=1}^6$ are nonzero, including one from each pair. Noting that $N = \text{area}(S)$, and that

$$|\Theta| \leq \text{area}(S) \leq 6|\Theta|,$$

this proves the result with a constant of the form $K_\mu = p_i/3$ for some i such that $p_i \neq 0$. \square

6.7 Proof of the patching theorem

We now give the proof of the patching theorem (Theorem 6.3.5), as described in Section 6.3. We refer throughout to the corresponding figures from the outline there.

Proof of Theorem 6.3.5. Since $\tau_1, \tau_2 \in \Omega$, A_n is balanced for all n . Thus by Corollary 6.2.6, if A_n is not tileable, it has a minimal counterexample U , i.e. a set with

$$\text{imbalance}(U) = \text{even}(U) - \text{odd}(U) > 0,$$

despite U having only odd cubes on its interior boundary. Let S denote the interior boundary surface of U . Since A_n is tileable with just the τ_1 boundary condition (resp. with just the τ_2 boundary condition), S must connect the inner and outer boundaries of A_n . Thus by Lemma 6.5.3, we can find $l_+ \in ((1 - \epsilon^{1/2})n, n)$ and $l_- \in ((1 - \delta)n, (1 - \delta + \epsilon^{1/2})n)$ such that for $l = l_+$ or $l = l_-$,

$$\text{Cov}_\epsilon(S \cap \partial B_l) \leq c\epsilon^{1/2}n^2 \quad (35)$$

$$\text{length}(\partial(S \cap \partial B_l)) \leq c'\epsilon^{-1/2}n \quad (36)$$

where c, c' are constants. We define the ‘‘middle region’’ $A_{\text{mid}} = (B_{l_+} \setminus B_{l_-})$, see Figure 25. Then we let

$$U' = U \cap A_{\text{mid}}.$$

See Figure 26. Let μ be an ergodic measure of mean current s (this exists by Lemma 6.6.1). Let Θ be the collection of odd cubes adjacent to $S' = S \cap U'$. Note that

$$\text{area}(S') \leq |\Theta| \leq 6\text{area}(S')$$

By Lemma 6.5.2, $|\Theta| \geq \text{area}(S') \geq c'_2 n^2$, where $c'_2 \sim (\delta - 2\epsilon^{1/2})^2$. Thus by Lemma 6.6.7 and Equation (36), there is a constant $K_\mu \in (0, 1)$ such that

$$\mathbb{E}_\mu[|\text{flux}(v_\tau, S')|] \geq K_\mu c'_2 n^2 + O(\epsilon^{-1/2}n).$$

By Theorem 6.6.3, for any coordinate plane P , a tiling τ sampled from μ is ϵ -nearly-constant on P with value s with probability approaching 1 as n goes to ∞ . Therefore we can sample a tiling τ from μ which is ϵ -nearly-constant with value s on A_{mid} for n large enough and satisfies

$$|\text{flux}(v_\tau, S')| \geq K_\mu c'_2 n^2 + O(\epsilon^{-1/2}n). \quad (37)$$

We fix this choice of τ for the rest of the proof. Define U_τ to be the region covered by the tiles from τ which intersect U' , see Figure 27. Since U_τ is tileable,

$$\text{imbalance}(U_\tau) = \text{even}(U_\tau) - \text{odd}(U_\tau) = 0.$$

We next define a new region U'' , which is U_τ minus the even cubes adjacent to S' (see Figure 27 again). Note that the region $U_\tau \setminus U''$ consists of only even cubes. By Equation (37),

$$|U_\tau \setminus U''| \geq K_\mu c'_2 n^2 + O(\epsilon^{-1/2} n).$$

Therefore

$$\text{imbalance}(U'') \leq -K_\mu c'_2 n^2 + O(\epsilon^{-1/2} n).$$

It remains to show that $\text{imbalance}(U)$ is very close to $\text{imbalance}(U'')$, and this is where we use the ϵ -nearly-constant condition. We split $A_n \setminus A_{\text{mid}}$ into two regions. First we define the “thin region” A_{thin} to be the union of columns parallel to one of the coordinate directions which connect ∂A_n and ∂A_{mid} . See Figures 25 and 28. The complement of $A_{\text{mid}} \cup A_{\text{thin}}$ we call the “corner region” A_{corner} , and consists of a neighborhood of the edges of the outer boundary cube of ∂A_n , and the inner boundary cube of ∂A_{mid} . We note that

$$\text{area}(A_{\text{corner}} \cap \partial A_n) \leq 24\epsilon^{1/2} n^2.$$

The constant factor comes from the fact that the cube has 12 edges. Therefore

$$\text{imbalance}(U \cap A_{\text{corner}}) \leq 24\epsilon^{1/2} n^2.$$

We define $U_{\text{shadow}} = (U \setminus U'') \cap A_{\text{thin}}$. We bound the imbalance of U_{shadow} column-by-column, where each column C consists of a straight line path of single cubes from ∂A_{mid} to ∂A_n . Recall U_{shadow} is defined with boundary condition τ on ∂A_{mid} and boundary condition τ_1 or τ_2 on ∂A_n . We also note that for any column C , since S is monochromatic black, $\text{imbalance}(U \cap C) \leq +1$.

We cut ∂A_{mid} into $(\epsilon n) \times (\epsilon n)$ patches α . For each α , there is a corresponding patch β on ∂A_n matched to α by columns. For any patch $\alpha \subset \partial A_{\text{mid}}$, since τ is sampled from an ergodic measure, it is ϵ -nearly-constant with value s on ∂A_{mid} (Corollary 6.6.5). Thus we have that

$$\text{flux}(v_\tau, \alpha) = \frac{1}{2} \langle s, \xi_\alpha \rangle (\epsilon n)^2 + o(\epsilon^2 n^2). \quad (38)$$

Let v_* be equal to v_{τ_1} on the outer boundary of ∂A_n and v_{τ_2} on the inner boundary. For $\beta \subset \partial A_n$ the patch connected by a column to α , since τ_1, τ_2 are ϵ -nearly-constant with value s ,

$$\text{flux}(v_*, \beta) = \frac{1}{2} \langle s, \xi_\alpha \rangle (\epsilon n)^2 + o(\epsilon^2 n^2). \quad (39)$$

(Note that $\xi_\alpha = \xi_\beta$.) For a patch $\alpha \subset \partial A_{\text{mid}}$, let $C(\alpha)$ be the union of columns incident to α . The set U_{shadow} is covered by these column sets, so it remains to control imbalance of $U_{\text{shadow}} \cap C(\alpha)$ for each patch α .

By Equation (35), at most $c\epsilon^{1/2} n^2$ of the area of ∂A_{mid} is in patches α which intersect U and U^c . The total imbalance contribution from these is bounded by $c\epsilon^{1/2} n^2$. If $\alpha \subset U^c$ but $C(\alpha)$ still intersects U , then all the columns in $C(\alpha)$ have at least one end on S , and the imbalance in $C(\alpha) \cap U_{\text{shadow}}$ is at most 0.

Now we look at the cases where $\alpha \subset U$. If $C(\alpha) \subset U$, then the total imbalance in $C(\alpha) \cap U$ is $o(\epsilon^2 n^2)$ by Equations (38) and (39). If $\alpha \subset U$ but $C(\alpha) \not\subset U$, then some of the columns starting from α hit S before hitting β . This means they end on an odd cube. Extending the column all the way to β would only make the imbalance larger, however then the imbalance in

$C(\alpha) \cap U$ is bounded above by $o(\epsilon^2 n^2)$ by Equations (38) and (39). The number of columns is bounded by a constant independent of n times ϵ^{-2} , hence in total the imbalance in U_{shadow} is bounded by $c\epsilon^{1/2}n^2 + o(n^2)$.

Putting everything together, we have that

$$\begin{aligned} \text{imbalance}(U) &\leq \text{imbalance}(U'') + \text{imbalance}(U_{\text{shadow}}) + \text{imbalance}(U \cap A_{\text{corner}}) \\ &\leq -K_\mu c'_2 n^2 + 24\epsilon^{1/2}n^2 + c\epsilon^{1/2}n^2 + O(\epsilon^{-1/2}n) + o(n^2). \end{aligned}$$

Recall that $c'_2 \sim (\delta - 2\epsilon^{1/2})^2$, so it gets larger as ϵ gets smaller. Fixing ϵ small enough as a function of the constants, for n large enough $\text{imbalance}(U) \leq 0$ and hence U is not a counterexample. Therefore by Proposition 6.2.6, for n large enough, A_n is tileable with boundary conditions τ_1, τ_2 . \square

6.8 Corollaries for ergodic Gibbs measures

Corollary 6.8.1. *If μ_1, μ_2 are EGMs of the same mean current $s \in \text{Int}(\mathcal{O})$, then $h(\mu_1) = h(\mu_2)$.*

Remark 6.8.2. The proof of the patching theorem (Theorem 6.3.5) uses that $s \notin \partial\mathcal{O}$ since this is a condition of Lemma 6.6.7. This corollary shows that $s \notin \partial\mathcal{O}$ is a necessary condition and not just an artifact of the proof, since if $s \in \partial\mathcal{O}$ then not all ergodic Gibbs measures of mean current s have the same specific entropy (see Proposition 4.2.10).

Proof. Fix $\delta > 0$, and let $B_n = [-n, n]^3$ and $A_n = B_n \setminus B_{(1-\delta)n}$. By the patching theorem (Theorem 6.3.5) with outer boundary condition sampled from μ_1 and inner boundary condition sampled from μ_2 , we get that

$$h(\mu_2) \leq (1 + O(\delta))h(\mu_1).$$

Switching them, we find that

$$h(\mu_1) \leq (1 + O(\delta))h(\mu_2).$$

Therefore $h(\mu_1) = h(\mu_2)$. \square

Another useful result comes from applying patching to a sequence of ϵ -nearly-constant tilings and a sample from an EGM. This relates the number of tilings of a region with fixed ϵ -nearly-constant boundary conditions to the specific entropy of an EGM. This serves as a lemma in the proof of the lower bound in the large deviation principle (Theorem 8.1.10). Recall that Ω denotes the set of dimer tilings of \mathbb{Z}^3 .

Proposition 6.8.3. *Fix $\delta > 0$, $s \in \text{Int}(\mathcal{O})$, $\epsilon > 0$ small enough, $B_n = [-n, n]^3$, and let $A_n = B_n \setminus B_{(1-\delta)n}$. Suppose that $(\tau_n)_{n \geq 1} \subset \Omega$ is such that τ_n is ϵ -nearly-constant on ∂B_n with value s for n large enough. Let π_n be the uniform probability measure on tilings σ of B_n such that $\sigma|_{\partial B_n} = \tau_n$. Then for any EGM μ of mean current s and n large enough,*

$$|B_n|^{-1}H(\pi_n) \geq h(\mu)(1 + O(\delta)).$$

Remark 6.8.4. Since π_n is a uniform measure, $H(\pi_n) = \log Z_n$, where Z_n is the partition function of π_n .

Proof. We apply the patching theorem (Theorem 6.3.5) to $A_n = B_n \setminus B_{(1-\delta)n}$ with τ_n on the outer boundary and a sample from μ on the inner boundary. For n large enough, patching is possible with μ -probability $(1 - \epsilon)$ on an annulus of width δ .

For $\Lambda \subset \mathbb{Z}^3$, let $\Omega(\Lambda)$ denote the free-boundary dimer tilings of Λ . For $\sigma \in \Omega(\Lambda)$, recall that $X(\sigma)$ is the set of *extensions* of σ , i.e.

$$X(\sigma) = \{\tilde{\sigma} \in \Omega : \tilde{\sigma}|_{\Lambda} = \sigma\}.$$

Then we compute

$$\begin{aligned} H_{B_{(1-\delta)n}}(\mu) &= - \sum_{\sigma \in \Omega(B_{(1-\delta)n})} \mu(X(\sigma)) \log \mu(X(\sigma)) \\ &= - \sum_{\substack{\sigma \in \Omega(B_{(1-\delta)n}) \\ \sigma, \tau_n \text{ patchable}}} \mu(X(\sigma)) \log \mu(X(\sigma)) + O(\epsilon \log \epsilon) \end{aligned}$$

Let U_n denote the uniform probability measure on

$$\{\sigma \in \Omega(B_{(1-\delta)n}) : \sigma, \tau_n \text{ patchable}\}.$$

Since uniform measure maximizes entropy,

$$H_{B_{(1-\delta)n}}(\mu) \leq H(U_n) + O(\epsilon \log \epsilon) \leq H(\pi_n) + O(\epsilon \log \epsilon).$$

Thus for n large enough such that the patching theorem applies for τ_n and a sample from μ on $A_n = B_n \setminus B_{(1-\delta)n}$, with probability $(1 - \epsilon)$ we have that

$$|B_n|^{-1} H(\pi_n) \geq (1 + O(\delta)) |B_{(1-\delta)n}|^{-1} H_{B_{(1-\delta)n}}(\mu) \geq (1 + O(\delta)) h(\mu).$$

□

7 Properties of entropy

In this section we prove results about the entropy functions ent and Ent introduced in Section 2.3. Recall that Ω is the set of dimer tilings of \mathbb{Z}^3 , and that \mathcal{P} denotes the space of $\mathbb{Z}^3_{\text{even}}$ -invariant probability measures on Ω . Further recall that for any $s \in \mathcal{O}$, we define $\mathcal{P}^s \subset \mathcal{P}$ to be the set of measures which also have mean current s , $\mathcal{P}_e \subset \mathcal{P}$ to be the set of ergodic measures, and \mathcal{P}_e^s to be the set of ergodic measures with mean current s . The mean-current entropy function $\text{ent} : \mathcal{O} \rightarrow [0, \infty)$ is defined by

$$\text{ent}(s) = \sup_{\mu \in \mathcal{P}^s} h(\mu),$$

where $h(\cdot)$ denotes specific entropy (see Section 2.3).

We saw in Section 4 that ent is equal to ent_{loz} when restricted to any face of $\partial\mathcal{O}$ (Theorem 4.2.7). In particular this implies that ent is strictly concave when restricted to any face of $\partial\mathcal{O}$ (Corollary 4.2.8). The main result of this section is that ent is strictly concave on all of $\mathcal{O} \setminus \mathcal{E}$, where \mathcal{E} denotes the edges of $\partial\mathcal{O}$ (Theorem 7.5.1).

In Section 7.1 we show that the supremum of $\{h(\mu) : \mu \in \mathcal{P}^s\}$ is realized by a Gibbs measure for all $s \in \mathcal{O}$ (Theorem 7.1.2). It is a classical result going back to Lanford and Ruelle [LR69]

that entropy maximizers in \mathcal{P} are Gibbs measures. We extend this to show that the entropy maximizer in \mathcal{P}^s , where the mean current is fixed, is also a Gibbs measure. In Section 7.2 we show using elementary methods that ent is concave and continuous on \mathcal{O} .

The proof that ent is strict concave on $\mathcal{O} \setminus \mathcal{E}$ (Theorem 7.5.1) requires some new tools and ideas. We use a version of a technique called *cluster swapping* used in [She05, Chapter 8], which we call *chain swapping*. As set up for the proof of strict concavity in Section 7.5 we prove some preliminary results about flows in the double dimer model (Section 7.3) and introduce the chain swapping technique (Section 7.4). Combining chain swapping with the results for ent on $\partial\mathcal{O}$ from Section 4 we show that ent is strictly concave on $\mathcal{O} \setminus \mathcal{E}$.

Strict concavity has a number of important consequences. We saw in Corollary 6.8.1 that if μ_1, μ_2 are EGMs with mean current $s \in \text{Int}(\mathcal{O})$, then $h(\mu_1) = h(\mu_2)$. Combining this with strict concavity, we show that if $s \in \text{Int}(\mathcal{O})$ and $\mu \in \mathcal{P}^s$, then $h(\mu) = \text{ent}(s)$ if and only if μ is an EGM or weighted average of EGMs all of mean current s (Theorem 7.5.2), and therefore that there exists an EGM of every mean-current $s \in \mathcal{O}$ (Corollary 7.5.4). Heuristically, the main goal of this section is to show that the mean current s captures broad statistics of dimer tilings sampled from $\mu \in \mathcal{P}_e^s$ when $s \in \text{Int}(\mathcal{O})$.

Finally in Section 7.6, we leverage properties of ent to study the entropy functional $\text{Ent} : AF(R) \rightarrow [0, \infty)$ on asymptotic flows given by

$$\text{Ent}(f) = \frac{1}{\text{Vol}(R)} \int_R \text{ent}(f(x)) \, dx.$$

The rate function for the large deviation principle in Section 8 will be $-\text{Ent}$ (up to an additive constant). Using the properties of ent from earlier in the section, we show that Ent is upper semicontinuous in the Wasserstein topology (Proposition 7.6.3) and strictly concave on the subspace of flows which never take values in the edges $\mathcal{E} \subset \partial\mathcal{O}$ (Corollary 7.6.1). After that, we adapt an argument of V. Gorin [Gor21] to show that there is a unique Ent maximizer in $AF(R, b)$ for any boundary asymptotic flow b under the mild condition that (R, b) is *semi-flexible* (Definition 7.6.8, Theorem 7.6.10).

7.1 Entropy maximizers of a given mean current are Gibbs measures

We first study the maximizers of specific entropy $h(\cdot)$ in \mathcal{P}^s for $s \in \mathcal{O}$ fixed. It is straightforward to show that there exists a measure $\mu \in \mathcal{P}^s$ which achieves $\sup\{h(\mu) : \mu \in \mathcal{P}^s\}$ for any $s \in \mathcal{O}$.

Lemma 7.1.1. *Let $s \in \mathcal{O}$. There exists $\mu \in \mathcal{P}^s$ such that $h(\mu) = \text{ent}(s)$.*

Proof. The space \mathcal{P}^s is compact with respect to the weak star topology. Since h is an upper semicontinuous function of the measure [Kel98, Theorem 4.2.4] it must achieve its maxima in \mathcal{P}^s . \square

Theorem 7.1.2. *Fix $s \in \mathcal{O}$. If $\mu \in \mathcal{P}^s$ has $h(\mu) = \text{ent}(s)$ then μ is a Gibbs measure.*

If the mean current is not fixed, then this is a standard result originally shown by Landford and Ruelle in [LR69]. The main idea of the proof is a variational argument which says that if a measure μ is not Gibbs, then there is a “perturbation” of μ which has more entropy. We need to show that this perturbation does not change the mean current, and this is the

purpose of Lemma 7.1.3. After that, our proof of Theorem 7.1.2 is inspired by the exposition in [BS94].

To show the mean current does not change, we use double dimers to compare the mean currents of the two measures. Double dimers will be a tool throughout Section 7. There is a natural action of the group $\mathbb{Z}_{\text{even}}^3$ on the product $\Omega \times \Omega$ acting coordinate wise. Superimposing two dimer tilings τ_1 and τ_2 gives us a *double dimer configuration* (τ_1, τ_2) . The union $\tau_1 \cup \tau_2$ consists of finite cycles, infinite paths and isolated double edges. Each cycle or infinite path in (τ_1, τ_2) is *oriented* in a way that agrees with the direction of the τ_1 flow (for edges in τ_1) or opposite the direction of the τ_2 flow (for edges in τ_2).

Lemma 7.1.3. *Let \mathbf{m} be a $\mathbb{Z}_{\text{even}}^3$ -invariant measure on $\Omega \times \Omega$ such that for \mathbf{m} almost every (τ_1, τ_2) , the union $\tau_1 \cup \tau_2$ does not contain infinite paths. Then*

$$s(\pi_1(\mathbf{m})) = s(\pi_2(\mathbf{m}))$$

where $\pi_i : \Omega \times \Omega \rightarrow \Omega$ is projection onto the i^{th} coordinate for $i = 1, 2$.

Remark 7.1.4. We remark that non-existence of infinite paths in a sample from a coupling like this has been used in other related but different ways in statistical mechanics, e.g. to show that two Gibbs measures are the same if there are no infinite paths in a sample from the coupling, or to compute covariances. See for example [vdB93], [vdBS94].

Proof. Since the mean current is an affine function of measure, by the ergodic decomposition theorem it is sufficient to prove this lemma for \mathbf{m} ergodic. For the rest of the proof we assume that \mathbf{m} is an ergodic measure.

Recall that v_τ is the *pretiling flow* of τ (for the definition see Equation (4) in Section 2.1), and let (τ_1, τ_2) be a sample from \mathbf{m} . By assumption all paths $\gamma \subset (\tau_1, \tau_2)$ are finite loops or double edges (which are loops with just one edge from each tiling). Let $E(\gamma)$ denote the edges along γ oriented from even to odd. (If $e \in E(\gamma)$, $-e$ is e with orientation reversed.) With a slight abuse of notation, we also view e as a vector oriented from even to odd. Since γ is a loop,

$$\sum_{e \in E(\gamma)} v_{\tau_1}(e)e = \sum_{e \in E(\gamma)} v_{\tau_2}(e)e. \quad (40)$$

Given $x \in \mathbb{Z}^3$ let γ_x denote the loop in (τ_1, τ_2) containing x . We denote the number of edges in a loop γ by $\text{length}(\gamma)$. For any $\epsilon > 0$, there exists k such that

$$\mathbf{m}(\text{length}(\gamma_0) > k) < \epsilon.$$

Let $B_n := [1, n]^3$. By the mean ergodic theorem, there is n large enough such that with \mathbf{m} -probability $1 - \epsilon$, the double dimer configuration (τ_1, τ_2) satisfies the following.

1. We have

$$|\{x \in B_n : \text{length}(\gamma_x) > k\}| < 2\epsilon n^3. \quad (41)$$

2. Let $E(B_n)$ denote the edges in B_n oriented from even to odd. For $i = 1, 2$ and any unit coordinate vector η ,

$$\left| \langle s(\pi_i(\mathbf{m})), \eta \rangle - \frac{2}{n^3} \sum_{e \in E(B_n)} \langle v_{\tau_i}(e)e, \eta \rangle \right| < \epsilon. \quad (42)$$

Let $C_n := \{x \in B_n : \gamma_x \subset B_n\}$ denote the \mathbf{m} -random set of points on loops in (τ_1, τ_2) contained in B_n . By Equation (41), with \mathbf{m} -probability $1 - \epsilon$,

$$|C_n| \geq (n - k)^3 - 2\epsilon n^3. \quad (43)$$

Clearly $\{\gamma_x : x \in C_n\}$ is a union of loops. Let $E(C_n)$ denote the edges of loops in this collection oriented from even to odd. By Equation (40),

$$\sum_{e \in E(C_n)} \left(v_{\tau_1}(e)e - v_{\tau_2}(e)e \right) = 0. \quad (44)$$

There are at most $\frac{n^3 - |C_n|}{2}$ tiles in τ_i in $E(B_n) \setminus E(C_n)$ for $i = 1, 2$. Therefore by Equation (43), with \mathbf{m} probability $1 - 2\epsilon$

$$\left| \frac{2}{n^3} \sum_{e \in E(B_n)} v_{\tau_i}(e)e - \frac{2}{n^3} \sum_{e \in E(C_n)} v_{\tau_i}(e)e \right| \leq \frac{n^3 - (n - k)^3 + 2\epsilon n^3}{n^3} = 1 + 2\epsilon - \frac{(n - k)^3}{n^3}. \quad (45)$$

Combining Equations (42), (44), (45) gives that for any unit coordinate vectors η ,

$$\left| \langle s(\pi_1(\mathbf{m})), \eta \rangle - \langle s(\pi_2(\mathbf{m})), \eta \rangle \right| < 6\epsilon + \frac{2n^3 - 2(n - k)^3}{n^3}. \quad (46)$$

Taking $n \rightarrow \infty$ and then $\epsilon \rightarrow 0$ completes the proof. □

To prove Theorem 7.1.2, we mimic the perturbative argument of [BS94, Proposition 1.19], applying Lemma 7.1.3 to show that this does not change the mean current.

Proof of Theorem 7.1.2. It suffices to show that if $\mu \in \mathcal{P}^s$ is not a Gibbs measure, then there exists $\nu \in \mathcal{P}^s$ such that $h(\nu) > h(\mu)$. Under the assumption that μ is not a Gibbs measure, there exists a finite set $R \subset \mathbb{Z}^3$ and a positive measure set $\Omega' \subset \Omega$ such that for all $\tau \in \Omega'$, the conditional distribution of possible extensions of $\tau|_{\mathbb{Z}^3 \setminus R}$ to a tiling of \mathbb{Z}^3 is not uniform. We can assume using stationarity that R is contained in the positive quadrant. Let $n \in \mathbb{N}$ be such that $R \subset [1, n - 1]^3$. Since the number of tilings of R depends only on the tiling restricted to $S := [0, n]^3 \setminus R$ there exists a tiling τ_0 in the support of μ such that the conditional distribution on the possible extensions of $\tau_0|_S$ to R is not uniform. Since entropy is maximized by the uniform measure we have that there is a $\delta > 0$ such that

$$H(\text{uniform distribution on extensions of } \tau_0|_S \text{ to } R) - H(\mu|_R \text{ conditioned on } \tau_0|_S) > \delta.$$

We now construct a modification of μ and show that it has the same mean current but more entropy. For this take a sample τ from μ . Let n be an odd integer and divide \mathbb{Z}^3 into translates of $B = [0, n]^3$ by $(n + 1)\mathbb{Z}^3$. For each such translated box B , resample τ in B conditioned on $\tau|_{\partial^\circ B}$, where $\partial^\circ B = B \setminus [1, n - 1]^3 \subset B$ is the inner boundary of B . This gives us a new measure ν on Ω , which is invariant with respect to the $(n + 1)\mathbb{Z}^3$ subaction. By averaging ν with respect to translations by elements of $[0, n]^3 \cap \mathbb{Z}_{\text{even}}^3$ we get a $\mathbb{Z}_{\text{even}}^3$ -invariant measure which we denote ν' .

Let \mathbf{m} be a measure on $\Omega \times \Omega$ which is a coupling of μ and ν' , where the sample from ν' is derived by the construction above from the μ sample. If (τ_1, τ_2) is sampled from \mathbf{m} , τ_1 and τ_2

differ only on the interiors of copies of B . Therefore (τ_1, τ_2) has no infinite paths \mathfrak{m} -a.s., so by Lemma 7.1.3,

$$s(\nu') = s(\mu).$$

On the other hand by the ergodic theorem, there exists $\epsilon > 0$ for which there is a $(n+1)\mathbb{Z}^3$ -invariant set $A \subset \Omega$ with the following properties:

1. $\mu(A) > \epsilon$.
2. For all $\tau \in A$, $\tau_0|_S$ appears in translated boxes B with density greater than ϵ .

Therefore

$$h(\nu') - h(\mu) > \frac{1}{(n+1)^3} \epsilon^2 \delta.$$

□

The proof of Theorem 7.1.2 also has a useful consequence for the double dimer model. Recall the maps $\pi_1, \pi_2 : \Omega \times \Omega \rightarrow \Omega$ given by $\pi_i(\tau_1, \tau_2) = \tau_i$. Let \mathcal{P}^{s_1, s_2} be the space of invariant probability measures μ on $\Omega \times \Omega$ such that $\pi_i(\mu) \in \mathcal{P}^{s_i}$ for $i = 1, 2$.

Corollary 7.1.5. *Let $s_1, s_2 \in \mathcal{O}$. Then*

$$\sup_{\mu \in \mathcal{P}^{s_1, s_2}} h(\mu) = \text{ent}(s_1) + \text{ent}(s_2).$$

Further, the measures $\mu \in \mathcal{P}^{s_1, s_2}$ which maximize specific entropy on \mathcal{P}^{s_1, s_2} are Gibbs measures on $\Omega \times \Omega$ and satisfy $h(\pi_i(\mu)) = \text{ent}(s_i)$ for $i = 1, 2$.

Proof. For any measure $\mu \in \mathcal{P}^{s_1, s_2}$ we have that

$$h(\mu) \leq h(\pi_1(\mu)) + h(\pi_2(\mu)) \leq \text{ent}(s_1) + \text{ent}(s_2). \quad (47)$$

For $i = 1, 2$, let $\nu_i \in \mathcal{P}^{s_i}$ be such that $h(\nu_i) = \text{ent}(s_i)$. The product measure $\nu = \nu_1 \times \nu_2$ has $h(\nu) = \text{ent}(s_1) + \text{ent}(s_2)$, so by Equation (47) ν maximizes specific entropy among measures in \mathcal{P}^{s_1, s_2} . Therefore if μ is a maximizer it must be in the equality case in Equation (47), which implies that $h(\pi_1(\mu)) = \text{ent}(s_1)$ and $h(\pi_2(\mu)) = \text{ent}(s_2)$.

Finally the proof that the entropy maximizer must be a Gibbs measure is exactly the same as the proof of Theorem 7.1.2; if a measure in \mathcal{P}^{s_1, s_2} is not a Gibbs measure then we can increase its entropy by locally modifying it. □

7.2 Basic properties of ent

In this section we prove some straightforward properties of the mean current entropy function ent. The only tools here are basic real analysis and properties of $h(\cdot)$.

Lemma 7.2.1. *ent is a concave function on \mathcal{O} .*

Proof. Fix $u, v \in \mathcal{O}$ and $\alpha \in (0, 1)$. By Lemma 7.1.1 we know that the entropy function $h(\cdot)$ on \mathcal{P}^u (resp. on \mathcal{P}^v) achieves a maximum say at μ (resp. at ν). Given this we have that $\alpha\mu + (1 - \alpha)\nu \in \mathcal{P}^{\alpha u + (1 - \alpha)v}$ and

$$h(\alpha\mu + (1 - \alpha)\nu) = \alpha h(\mu) + (1 - \alpha)h(\nu) = \alpha \text{ent}(u) + (1 - \alpha)\text{ent}(v).$$

Thus

$$\text{ent}(\alpha u + (1 - \alpha)v) \geq \alpha \text{ent}(u) + (1 - \alpha) \text{ent}(v)$$

which shows that ent is concave. \square

Lemma 7.2.2. *ent is an upper semi-continuous function on \mathcal{O} .*

Proof. Let $u_n \in \mathcal{O}$ be a sequence such that $u_n \rightarrow u$. By Lemma 7.1.1 there exist a measure μ_n maximizing the entropy function $h(\cdot)$ on \mathcal{P}^{u_n} . Since the mean current is a continuous function of the measure (Definition 2.2.2), we have that any subsequential limit μ of μ_n must lie in \mathcal{P}^u . Since $h(\cdot)$ is upper semicontinuous as a function on \mathcal{P} ,

$$\text{ent}(u) \geq h(\mu) \geq \limsup_{n \rightarrow \infty} h(\mu_n) = \limsup_{n \rightarrow \infty} \text{ent}(u_n),$$

which completes the proof. \square

We put these together to show that ent is continuous.

Lemma 7.2.3. *ent is a continuous function on \mathcal{O} .*

Proof. Fix $u \in \mathcal{O}$ and $\epsilon > 0$. Let $M := \sup_{v \in \mathcal{O}} \text{ent}(v)$; M is finite since \mathcal{O} is compact and ent is upper semicontinuous on \mathcal{O} (Lemma 7.2.2). Let $\|\cdot\|_1$ denote the L^1 norm. Again using Lemma 7.2.2, there exists $\delta_1 > 0$ such that if $v \in \mathcal{O}$ is such that $\|v - u\|_1 < \delta_1$, then

$$\text{ent}(u) - \text{ent}(v) > -\epsilon.$$

Choose $L > 1 + (2M/\epsilon)$ and $\delta_2 > 0$ small enough such that if $\|v - u\|_1 < \delta_2$ then $u + L(v - u) \in \mathcal{O}$ (note that this is possible even when $u \in \partial\mathcal{O}$ since \mathcal{O} is convex). By Lemma 7.2.1,

$$\frac{L - 1}{L} \text{ent}(u) + \frac{1}{L} \text{ent}(u + L(v - u)) \leq \text{ent}(v).$$

Rearranging the equation we get that

$$\text{ent}(u) - \text{ent}(v) \leq \frac{1}{L} \text{ent}(u) - \frac{1}{L} \text{ent}(u + L(v - u)) \leq 2M/L \leq \epsilon.$$

Taking $\delta < \delta_1 \delta_2$, if $\|v - u\|_1 < \delta$ then $|\text{ent}(u) - \text{ent}(v)| < \epsilon$, which proves that ent is continuous. \square

7.3 Flows for the double dimer model

To prove that ent is strictly concave we use a double dimer model construction called *chain swapping*, which is an operation on infinite paths in a double dimer configuration related to the cluster swapping technique used in [She05]. In this section we give some of the necessary background results for the double dimer model, and in the next section we explain what chain swapping is.

Here we look at $\mathbb{Z}_{\text{even}}^3$ -invariant couplings of $\mathbb{Z}_{\text{even}}^3$ -invariant measures $\mu_1, \mu_2 \in \mathcal{P}_e$ and study properties of the sample (τ_1, τ_2) . We distinguish between the pair of dimer tilings (τ_1, τ_2) and the union of tilings $\tau_1 \cup \tau_2$, where we forget the information of which tiling each edge $e \in \tau_1 \cup \tau_2$ belongs to. As we observed earlier, $\tau_1 \cup \tau_2$ is a set of isolated double edges, finite cycles, and infinite paths. We saw in Lemma 7.1.3 that if $\tau_1 \cup \tau_2$ consists of only finite cycles and double edges, then the marginals have the same mean current. This suggests that the infinite paths in a double dimer configuration carry a lot of information about the difference between the mean currents of the measures involved. The main results of this section are Proposition 7.3.2 and Corollary 7.3.3 which make this precise.

Recall from Section 2.1 that the *flow associated with a double dimer configuration* (τ_1, τ_2) is

$$f_{(\tau_1, \tau_2)} = v_{\tau_1} - v_{\tau_2},$$

where v_τ is the pretiling flow defined in Section 2.1, Equation (4). (Equivalently, $f_{(\tau_1, \tau_2)} = f_{\tau_1} - f_{\tau_2}$ where f_τ is the tiling flow, since the reference flows cancel.) Explicitly, for each edge e oriented from even to odd,

$$f_{(\tau_1, \tau_2)}(e) = \begin{cases} 1 & \text{if } e \in \tau_1 \setminus \tau_2 \\ -1 & \text{if } e \in \tau_2 \setminus \tau_1 \\ 0 & \text{if } e \in \tau_1 \cap \tau_2 \text{ or if } e \notin \tau_1 \cup \tau_2. \end{cases}$$

The vector field $f_{(\tau_1, \tau_2)}$ is divergence free, and its flow lines are the cycles and paths of the double dimer configuration $\tau_1 \cup \tau_2$. In particular, each $x \in \mathbb{Z}^3$ is in one of two cases:

1. $f_{(\tau_1, \tau_2)}$ is equal to 1 on exactly two edges e_1, e_2 incident to x , with one of the edges oriented into x and the other oriented out of x .
2. $f_{(\tau_1, \tau_2)}$ is zero on all edges e incident to x .

The set of vertices $x \in \mathbb{Z}^3$ in Case 2 is the collection of vertices covered by $\tau_1 \cap \tau_2$. In particular, we note that it is tileable by dimers.

Conversely, if a discrete vector field g satisfies these properties (i.e. all vertices are in Case 1 or Case 2, and the set of Case 2 vertices is tileable by dimers), then there exist tilings τ_1, τ_2 such that $g = f_{(\tau_1, \tau_2)}$. In fact we can explicitly construct the tilings from g . Given any tiling τ of the Case 2 vertices $\{v \in \mathbb{Z}^3 : g(e) = 0 \text{ for all } e \text{ incident to } v\}$, we define

$$\begin{aligned} \tau_1 &= \tau \cup \{e : g(e) = 1 \text{ where } e \text{ is an edge directed from an even to an odd vertex}\} \\ \tau_2 &= \tau \cup \{e : g(e) = -1 \text{ where } e \text{ is an edge directed from an even to an odd vertex}\}. \end{aligned}$$

From this we see that that the flow $f_{(\tau_1, \tau_2)}$ determines the double dimer configuration (τ_1, τ_2) up to the choice of tiling τ on the Case 2 vertices. On other other hand, the union of tilings $\tau_1 \cup \tau_2$ determines (τ_1, τ_2) on the set where $\tau_1 = \tau_2$, meaning it determines the tiling τ of the Case 2 vertices. Therefore together these are enough to recover (τ_1, τ_2) . In summary, we have shown:

Proposition 7.3.1. *The pair $(\tau_1 \cup \tau_2, f_{(\tau_1, \tau_2)})$ uniquely determines the double dimer configuration (τ_1, τ_2) and vice versa.*

Shifting along flow lines. We define a \mathbb{Z} -action on $\Omega \times \Omega$ by translating in the direction of the double dimer flow $f_{(\tau_1, \tau_2)}$. Given $(\tau_1, \tau_2) \in \Omega \times \Omega$, let $b_1 \in \tau_1$ be the edge incident to the origin, and suppose that $b_1 = (0, a_1)$, $a_1 \in \mathbb{Z}^3$ a neighbor of the origin. Following that, let $b_2 \in \tau_2$ be the edge incident to a_1 , and suppose that $b_2 = (a_1, a_1 + a_2)$, where $a_2 \in \mathbb{Z}^3$ a neighbor of the origin. These are the first two edges of a path in (τ_1, τ_2) . We define $\alpha_1(\tau_1, \tau_2)$ to be the directed vector a_1 and $\alpha_2(\tau_1, \tau_2)$ to be the directed vector a_2 . When the pair of tilings (τ_1, τ_2) is understood, we drop them from the notation. We then define the function $F : \Omega \times \Omega \rightarrow [-2, 2]^3$ by

$$F((\tau_1, \tau_2)) = \alpha_1 + \alpha_2.$$

We define F to be translation by two edges so that the parity of the even/odd vertices is preserved. This can be viewed as tracking the slope and speed of the flow $f_{(\tau_1, \tau_2)}$ (when there is a double edge at the origin in (τ_1, τ_2) , F is 0). Finally we define a transformation $T : \Omega \times \Omega \rightarrow \Omega \times \Omega$ given by translating the double dimer tiling by $\alpha_1 + \alpha_2$. If (τ_1, τ_2) has a double edge at the origin, then $T((\tau_1, \tau_2)) = (\tau_1, \tau_2)$. Otherwise, T shifts (τ_1, τ_2) along the path through the origin. The corresponding flow $T(f_{(\tau_1, \tau_2)})(e) = f_{(\tau_1, \tau_2)}(e + \alpha_1 + \alpha_2)$. If μ is a $\mathbb{Z}_{\text{even}}^3$ -invariant measure on $\Omega \times \Omega$, then it is also T -invariant. Thus by the ergodic theorem we have that for μ almost every $(\tau_1, \tau_2) \in \Omega \times \Omega$,

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N F(T^i((\tau_1, \tau_2))) =: F^*((\tau_1, \tau_2))$$

exists. Further, F^* is invariant under T and

$$\int_{\Omega \times \Omega} F^*((\tau_1, \tau_2)) \, d\mu((\tau_1, \tau_2)) = \int_{\Omega \times \Omega} F((\tau_1, \tau_2)) \, d\mu((\tau_1, \tau_2)).$$

By construction, F^* measures the slope or asymptotic direction of the path γ_0 through the origin in (τ_1, τ_2) . We call F^* the *slope function*. If γ_0 is a double edge or finite cycle, then F^* is 0. If γ_0 is an infinite path, then it can have nonzero slope.

For any infinite path $\ell \subset (\tau_1, \tau_2)$ we can compute its slope by translating so that ℓ goes through the origin. We say that ℓ has *nonzero slope* if $F^*((\tilde{\tau}_1, \tilde{\tau}_2)) \neq 0$, where $(\tilde{\tau}_1, \tilde{\tau}_2)$ is a translation of (τ_1, τ_2) so that ℓ contains the origin (this is well-defined since F^* is T -invariant). With this we can prove the main result of this section.

Proposition 7.3.2. *Let μ be a measure on $\Omega \times \Omega$ which is a $\mathbb{Z}_{\text{even}}^3$ -invariant coupling of $\mathbb{Z}_{\text{even}}^3$ -invariant measures μ_1 and μ_2 on Ω . Then*

$$\int_{\Omega \times \Omega} F^*((\tau_1, \tau_2)) \, d\mu((\tau_1, \tau_2)) = s(\mu_1) - s(\mu_2).$$

Proof. Since μ is T -invariant,

$$\int_{\Omega \times \Omega} F^*((\tau_1, \tau_2)) \, d\mu((\tau_1, \tau_2)) = \int_{\Omega \times \Omega} (\alpha_1(\tau_1, \tau_2) + \alpha_2(\tau_1, \tau_2)) \, d\mu((\tau_1, \tau_2)).$$

Since $\alpha_1(\tau_1, \tau_2)$ is the vector along the edge adjacent to 0 in τ_1 pointing away from it, it depends only on μ_1 . Hence

$$\int_{\Omega \times \Omega} \alpha_1(\tau_1, \tau_2) \, d\mu((\tau_1, \tau_2)) = \int_{\Omega} \alpha_1(\tau_1, \tau_2) \, d\mu_1(\tau_1) = s(\mu_1).$$

The vector $\alpha_2(\tau_1, \tau_2)$ is defined similarly, but first we have to sum over the possible values of α_1 .

$$\int_{\Omega \times \Omega} \alpha_2(\tau_1, \tau_2) d\mu((\tau_1, \tau_2)) = \sum_{a_1, a_2 \in \star} a_2 \mu((0, a_1) \in \tau_1, (a_1, a_1 + a_2) \in \tau_2)$$

where \star is the six neighbors of the origin. By the $\mathbb{Z}_{\text{even}}^3$ -invariance of μ we get that

$$\begin{aligned} \int_{\Omega \times \Omega} \alpha_2(\tau_1, \tau_2) d\mu((\tau_1, \tau_2)) &= \sum_{a_1, a_2 \in \star} a_2 \mu((-a_1 - a_2, -a_2) \in \tau_1, (-a_2, 0) \in \tau_2) \\ &= \sum_{a_2 \in \star} a_2 \mu_2((-a_2, 0) \in \tau_2) = -s(\mu_2). \end{aligned}$$

This completes the proof. \square

As a corollary, we show that the mean current difference of a pair of measures (μ_1, μ_2) can be computed by looking only at the tiles on infinite paths of nonzero slope. As a consequence, note also that if $s(\mu_1) \neq s(\mu_2)$ then an invariant coupling must have order n^3 tiles along infinite paths of nonzero slope. Here recall that $s_0(\tau)$ denotes the tile direction at the origin in τ and that for a measure μ_1 on Ω , $s(\mu_1) = \mathbb{E}_\mu[s_0(\tau)]$.

Corollary 7.3.3. *Let μ be a measure on $\Omega \times \Omega$ which is a $\mathbb{Z}_{\text{even}}^3$ -invariant coupling of $\mathbb{Z}_{\text{even}}^3$ -invariant measures μ_1 and μ_2 on Ω . Let I_0 be the event that the origin is contained in an infinite path of nonzero slope in (τ_1, τ_2) , and let*

$$s(\mu_1, I_0) - s(\mu_2, I_0) = \mathbb{E}_\mu[(s_0(\tau_1) - s_0(\tau_2)) \mathbb{1}_{I_0}((\tau_1, \tau_2))].$$

Then

$$s(\mu_1) - s(\mu_2) = s(\mu_1, I_0) - s(\mu_2, I_0).$$

Proof. Note that

$$\int_{\Omega \times \Omega} F^*((\tau_1, \tau_2)) d\mu((\tau_1, \tau_2)) = \int_{\Omega \times \Omega} F^*((\tau_1, \tau_2)) \mathbb{1}_{\{F^* \neq 0\}}((\tau_1, \tau_2)) d\mu((\tau_1, \tau_2))$$

By Proposition 7.3.2, the left hand side is equal to $s(\mu_1) - s(\mu_2)$. Since $I_0 = \{F^* \neq 0\}$, the right hand side is equal to $s(\mu_1, I_0) - s(\mu_2, I_0)$. \square

Finally we observe that the number of infinite paths of nonzero slope in (τ_1, τ_2) that intersect two far away boxes is 0 with probability 1 as the distance between the boxes goes to ∞ . This serves as a lemma for Proposition 7.4.2.

Lemma 7.3.4. *Let μ be a $\mathbb{Z}_{\text{even}}^3$ -invariant probability measure on $\Omega \times \Omega$, and fix $m \in \mathbb{N}$. Given a sample (τ_1, τ_2) from μ and $x, y \in \mathbb{Z}_{\text{even}}^3$, let $L_{x,y}$ denote the number of infinite paths of nonzero slope in (τ_1, τ_2) which intersect both $x + [1, m^3]$ and $y + [1, m^3]$. Then*

$$\lim_{n \rightarrow \infty} \frac{1}{n^3} \sum_{x \in [1, n^3]} \mu(L_{0,x} = 0) = 1.$$

Proof. There are at most m^3 infinite paths of nonzero slope in (τ_1, τ_2) passing through $[1, m]^3$. On the other hand, for any infinite path ℓ with nonzero slope,

$$|\ell \cap [1, n]^3| = O(n).$$

Therefore each infinite path ℓ of nonzero slope intersecting $[1, m]^3$ intersects $x + [1, m]^3$ for at most $O(n)$ points $x \in [1, n]^3$. Since m^3 is a constant, this implies that the number of $x \in [1, n]^3$ such that $L_{0,x} \neq 0$ is also $O(n)$. We can rewrite

$$\frac{1}{n^3} \sum_{x \in [1, n^3]} \mu(L_{0,x} \neq 0) = \mathbb{E}_\mu \left(\frac{1}{n^3} \sum_{x \in [1, n^3]} \mathbb{1}_{(L_{0,x} \neq 0)}(\tau_1, \tau_2) \right).$$

By the dominated convergence theorem, the right hand side tends to 0 as $n \rightarrow \infty$. This completes the proof. \square

7.4 Chain swapping

We can now introduce the main tool of this section, namely *chain swapping*, which is an operation on double dimer configurations similar to the *cluster swapping* technique used in [She05, Chapter 8].

Let $(\tau_1, \tau_2) \in \Omega \times \Omega$ be a pair of dimer tilings. Corresponding to this are the collection of (unoriented) loops $\tau_1 \cup \tau_2$ and the double dimer flow $f_{(\tau_1, \tau_2)}$. The flow $f_{(\tau_1, \tau_2)}$ determines the orientation of each loop or infinite path in $\tau_1 \cup \tau_2$.

For a fixed $p \in (0, 1)$, from a random configuration (τ_1, τ_2) we define a new random pair (τ'_1, τ'_2) by “shifting” the tiles along each infinite path of nonzero slope $\ell \subset (\tau_1, \tau_2)$ with independent probability p . In terms of the flow $f_{(\tau_1, \tau_2)}$, shifting on the infinite path ℓ corresponds to flipping the sign of $f_{(\tau_1, \tau_2)}$ along ℓ . The new tilings τ'_1, τ'_2 have the following properties:

1. $\tau_1 \cup \tau_2 = \tau'_1 \cup \tau'_2$, i.e. they correspond to the same collection of double edges, finite loops, and infinite paths.
2. Let ℓ_1, ℓ_2, \dots be the infinite paths of nonzero slope in (τ_1, τ_2) . Independently for each i , either with probability $1 - p$ the tiles on ℓ were not swapped, in which case

$$\tau'_1 \cap \ell_i = \tau_1 \cap \ell_i, \quad \tau'_2 \cap \ell_i = \tau_2 \cap \ell_i$$

or with probability p the tiles were swapped, in which case

$$\tau'_1 \cap \ell_i = \tau_1 \cap \ell_i, \quad \tau'_2 \cap \ell_i = \tau_2 \cap \ell_i.$$

3. On the complement of the infinite paths with nonzero slope in (τ_1, τ_2) , τ'_1 is equal to τ_1 and τ'_2 is equal to τ_2 .

We call this procedure *chain swapping with probability p* . See Figure 29. Chain swapping transforms a measure μ on $\Omega \times \Omega$ into a new measure μ' on $\Omega \times \Omega$ which we call the *swapped measure*.

Remark 7.4.1. Note that we only swap on the infinite paths of *nonzero slope*. This is a technical point. We do not know if the “asymptotic independence” result of Lemma 7.3.4 holds for infinite paths of zero slope. However we need Lemma 7.3.4 to show that the swapped measure is still ergodic (Proposition 7.4.2).




Figure 29: An example of tilings τ_1, τ_2 , the loops in (τ_1, τ_2) , and chain swapped tilings τ'_1, τ'_2 .

For the rest of this section, we study whether or not certain properties (ergodicity, the Gibbs property) are preserved under chain swapping, and how certain quantities (entropy, mean current) transform under chain swapping.

The first result is that chain swapping preserves ergodicity.

Proposition 7.4.2. *If μ is a ergodic measure on $\Omega \times \Omega$ with respect to the $\mathbb{Z}_{\text{even}}^3$ action and μ' is obtained from μ by chain swapping with probability p , then μ' is also ergodic.*

Proof. In this proof, we have two different parameters n (parameterizing possible translations of boxes) and m (the size of the boxes). Let $B_n = [1, n]^3$ and $\text{even}(B_n) = B_n \cap \mathbb{Z}_{\text{even}}^3$. Let $B_m + x$ denote B_m translated by $x \in \mathbb{Z}_{\text{even}}^3$. It is enough to show that for any two double dimer patterns restricted to B_m , denoted Π_1, Π_2 ,

$$\begin{aligned} & \lim_{n \rightarrow \infty} \frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu' \left((\tau'_1, \tau'_2) |_{B_m} = \Pi_1, (\tau'_1, \tau'_2) |_{B_m+x} = \Pi_2 \right) \\ &= \mu' \left((\tau'_1, \tau'_2) |_{B_m} = \Pi_1 \right) \mu' \left((\tau'_1, \tau'_2) |_{B_m+x} = \Pi_2 \right). \end{aligned}$$

Define the random variable L_x to be the number of infinite paths of nonzero slope in (τ_1, τ_2) sampled from μ which intersect $B_m + x$. Similarly define $L_{0,x}$ to be the number of infinite paths of nonzero slope which intersect B_m and $B_m + x$. Since the collection of tiles on infinite paths of nonzero slope is the same for μ and μ' , the quantities L_0 and $L_{0,x}$ are preserved by chain swapping. Let (τ_1, τ_2) have law μ and (τ'_1, τ'_2) have law μ' . Finally let \mathbf{m} be the coupling of μ, μ' given by chain swapping. Then

$$\frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu' \left((\tau'_1, \tau'_2) |_{B_m} = \Pi_1, (\tau'_1, \tau'_2) |_{B_m+x} = \Pi_2 \right)$$

$$\begin{aligned}
&= \frac{1}{\text{even}(B_n)} \sum_{\substack{k_1, k_2, k_3 \geq 0 \\ \Sigma_1, \Sigma_2 \text{ double dimer} \\ \text{tilings of } B_m}} \sum_{x \in \text{even}(B_n)} \mathfrak{m} \left((\tau_1, \tau_2) \mid_{B_m = \Sigma_1}, (\tau_1, \tau_2) \mid_{B_m+x = \Sigma_2}, \right. \\
&\quad \left. (\tau'_1, \tau'_2) \mid_{B_m = \Pi_1}, (\tau'_1, \tau'_2) \mid_{B_m+x = \Pi_2}, \right. \\
&\quad \left. L_0 = k_1, L_x = k_2, L_{0,x} = k_3 \right).
\end{aligned}$$

For each infinite path of nonzero slope in (τ_1, τ_2) we have an independent probability p of reversing its direction. For any triple $l = (l_1, l_2, l_3)$ with $l_i \leq k_i$ for each i , we define the notation

$$q_{k,l} = p^{l_1+l_2-l_3} (1-p)^{k_1+k_2-k_3-l_1-l_2+l_3}.$$

This is the probability of switching (l_1, l_2, l_3) of the (k_1, k_2, k_3) paths. With this notation, for each $x \in B_n$, the x term in the sum above is equal to

$$\begin{aligned}
&\sum_{\substack{k_1, k_2, k_3 \geq 0 \\ \Sigma_1, \Sigma_2 \text{ double dimer} \\ \text{tilings of } B_m \text{ which can swap} \\ \Pi_1, \Pi_2 \text{ with } (l_1, l_2, l_3) \text{ swaps}}} \mu \left((\tau_1, \tau_2) \mid_{B_m = \Sigma_1}, (\tau_1, \tau_2) \mid_{B_m+x = \Sigma_2}, L_0 = k_1, L_x = k_2, L_{0,x} = k_3 \right) q_{k,l}.
\end{aligned} \tag{48}$$

For any $K > 0$,

$$\frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \left(k_3 = K \text{ term in Equation (48)} \right) \leq \frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu(L_{0,x} = K).$$

By Proposition 7.3.4, the right hand side goes to 0 as $n \rightarrow \infty$. Therefore in the limit as $n \rightarrow \infty$, it suffices to consider the terms where $k_3 = 0$ (corresponding to the set of lines hitting B_n and the set of infinite paths hitting $B_n + x$ being disjoint). Therefore

$$\begin{aligned}
&\frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu' \left((\tau'_1, \tau'_2) \mid_{B_m = \Pi_1}, (\tau'_1, \tau'_2) \mid_{B_m+x = \Pi_2} \right) \\
&= \sum_{\substack{k_1, k_2 \geq 0 \\ (\Sigma_1, \Sigma_2) \text{ double dimer on } B_m \text{ which can swap to } \Pi_1, \Pi_2 \\ \text{with } (l_1, l_2) \text{ swaps}}} \\
&\frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu \left((\tau_1, \tau_2) \mid_{B_m = \Sigma_1}, (\tau_1, \tau_2) \mid_{B_m+x = \Sigma_2}, L_0 = k_1, L_x = k_2 \right) r_{k,l} + o(1),
\end{aligned}$$

where $r_{k,l} = p^{l_1+l_2} (1-p)^{k_1+k_2-l_1-l_2}$ (i.e. $q_{k,l}$ when $k_3 = 0$). Since μ is ergodic, for each $\Sigma_1, \Sigma_2, k_1, k_2$,

$$\begin{aligned}
&\lim_{n \rightarrow \infty} \frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu \left((\tau_1, \tau_2) \mid_{B_m = \Sigma_1}, (\tau_1, \tau_2) \mid_{B_m+x = \Sigma_2}, L_0 = k_1, L_x = k_2 \right) r_{k,l} \\
&= \mu \left((\tau_1, \tau_2) \mid_{B_m = \Sigma_1}, L_0 = k_1 \right) \mu \left((\tau_1, \tau_2) \mid_{B_m+x = \Sigma_2}, L_x = k_2 \right) r_{k,l}
\end{aligned}$$

Therefore

$$\lim_{n \rightarrow \infty} \frac{1}{\text{even}(B_n)} \sum_{x \in \text{even}(B_n)} \mu' \left((\tau'_1, \tau'_2) \mid_{B_m = \Pi_1}, (\tau'_1, \tau'_2) \mid_{B_m+x = \Pi_2} \right)$$

$$\begin{aligned}
&= \sum_{\substack{k_1, k_2 \geq 0 \\ \Sigma_1, \Sigma_2 \text{ double dimer tilings of } B_m \\ \text{which can swap to } \Pi_1, \Pi_2 \\ \text{with } (l_1, l_2) \text{ swaps}}} \mu \left((\tau_1, \tau_2) \mid_{B_m} = \Sigma_1, L_0 = k_1 \right) \mu \left((\tau_1, \tau_2) \mid_{B_{m+x}} = \Sigma_2, L_x = k_2 \right) r_{k,l} \\
&= \mu' \left((\tau'_1, \tau'_2) \mid_{B_m} = \Pi_1 \right) \mu' \left((\tau'_1, \tau'_2) \mid_{B_{m+x}} = \Pi_2 \right).
\end{aligned}$$

□

We now see how chain swapping affects the entropy and mean current of the marginal distributions.

Proposition 7.4.3. *Let μ be a measure on $\Omega \times \Omega$ which is an ergodic coupling of ergodic measures $\mu_1 \in \mathcal{P}_e^{s_1}$ and $\mu_2 \in \mathcal{P}_e^{s_2}$. If μ' is the measure obtained from μ by chain swapping with probability $p \in (0, 1)$, then $h(\mu') = h(\mu)$.*

Proof. In this proof, for a stationary random field X we let $h(X)$ denote the specific entropy of the law of X . Let (τ_1, τ_2) be a sample from μ and (τ'_1, τ'_2) be obtained by chain swapping. By Proposition 7.3.1,

$$h(\mu') = h((\tau'_1, \tau'_2)) = h(\tau'_1 \cup \tau'_2) + h(f_{(\tau'_1, \tau'_2)} \mid \tau'_1 \cup \tau'_2).$$

Since chain swapping preserves the set of tiles, $\tau'_1 \cup \tau'_2 = \tau_1 \cup \tau_2$, and we have automatically that $h(\tau'_1 \cup \tau'_2) = h(\tau_1 \cup \tau_2)$. On the other hand note that

$$\begin{aligned}
h(f_{(\tau'_1, \tau'_2)}, f_{(\tau_1, \tau_2)} \mid \tau'_1 \cup \tau'_2) &= h(f_{(\tau_1, \tau_2)} \mid \tau_1 \cup \tau_2) + h(f_{(\tau'_1, \tau'_2)} \mid \tau_1 \cup \tau_2, f_{(\tau_1, \tau_2)}) \\
&= h(f_{(\tau'_1, \tau'_2)} \mid \tau_1 \cup \tau_2) + h(f_{(\tau_1, \tau_2)} \mid \tau_1 \cup \tau_2, f_{(\tau'_1, \tau'_2)})
\end{aligned}$$

Conditioned on $\tau_1 \cup \tau_2$ and $f_{(\tau_1, \tau_2)}$, the distribution of the flow $f_{(\tau'_1, \tau'_2)}$ is determined by independent random choices for the orientation of each infinite path of nonzero slope in $\tau_1 \cup \tau_2$. Let $B_n = [1, n]^3$, and let $\ell \subset \tau_1 \cup \tau_2$ be an infinite path of nonzero slope. If $\ell \cap B_n$ is nonempty, then the orientation of ℓ is determined by its direction when it intersects ∂B_n . Therefore there exists a constant $c > 0$ such that

$$h(f_{(\tau'_1, \tau'_2)} \mid \tau_1 \cup \tau_2, f_{(\tau_1, \tau_2)}) \leq \lim_{n \rightarrow \infty} \frac{|\partial B_n|}{|B_n|} \leq \lim_{n \rightarrow \infty} \frac{cn^2}{n^3} = 0.$$

We can analogously show that $h(f_{(\tau_1, \tau_2)} \mid \tau_1 \cup \tau_2, f_{(\tau'_1, \tau'_2)}) = 0$. Therefore

$$h(f_{(\tau'_1, \tau'_2)} \mid \tau'_1 \cup \tau'_2) = h(f_{(\tau_1, \tau_2)} \mid \tau_1 \cup \tau_2)$$

so

$$h(\mu') = h(\tau_1 \cup \tau_2) + h(f_{(\tau_1, \tau_2)} \mid \tau_1 \cup \tau_2) = h(\mu).$$

□

Proposition 7.4.4. *Let μ be a measure on $\Omega \times \Omega$ which is an ergodic coupling of ergodic measures μ_1, μ_2 with mean currents $s(\mu_1), s(\mu_2)$. If μ' is the measure obtained from μ by chain swapping with probability $p \in (0, 1)$, then the marginal measures $\mu'_1 = \pi_1(\mu')$ and $\mu'_2 = \pi_2(\mu')$ have mean currents*

$$\begin{aligned}
s(\mu'_1) &= (1-p)s(\mu_1) + ps(\mu_2) \\
s(\mu'_2) &= ps(\mu_1) + (1-p)s(\mu_2).
\end{aligned}$$

Proof. As in the previous section let I_0 be the event that the origin is contained on an infinite path of nonzero slope in (τ_1, τ_2) . By Corollary 7.3.3,

$$s(\mu_1) - s(\mu_2) = s(\mu_1, I_0) - s(\mu_2, I_0),$$

where $s(\mu, I_0)$ is shorthand for the mean current computed as an average over only tilings where the origin is along an infinite path of nonzero slope. Since chain swapping only changes tiles that are contained on infinite paths of nonzero slope,

$$\begin{aligned} s(\mu'_1) - s(\mu_1) &= s(\mu'_1, I_0) - s(\mu_1, I_0) \\ s(\mu'_2) - s(\mu_2) &= s(\mu'_2, I_0) - s(\mu_2, I_0). \end{aligned}$$

On the other hand, since each infinite path of nonzero slope is swapped with independent probability p ,

$$\begin{aligned} s(\mu'_1, I_0) &= (1-p)s(\mu_1, I_0) + ps(\mu_2, I_0) \\ s(\mu'_2, I_0) &= ps(\mu_1, I_0) + (1-p)s(\mu_2, I_0). \end{aligned}$$

Combining gives

$$s(\mu'_1) - s(\mu_1) = -ps(\mu_1, I_0) + ps(\mu_2, I_0) = -ps(\mu_1) + ps(\mu_2).$$

Therefore $s(\mu'_1) = (1-p)s(\mu_1) + ps(\mu_2)$. An analogous calculation gives the result for $s(\mu'_2)$. \square

Finally we will show that chain swapping does **not** preserve the Gibbs property. To do this, we need two technical lemmas about double dimer configurations. This result is more involved than the other chain swapping results, so for simplicity we only prove this in the $p = 1/2$ case.

Let μ be a measure on $\Omega \times \Omega$ which is an ergodic coupling of ergodic measures μ_1, μ_2 on Ω such that $s(\mu_1) \neq s(\mu_2)$. Let P be a plane with normal vector ξ such that $\langle s(\mu_1) - s(\mu_2), \xi \rangle \neq 0$. Given a sample (τ_1, τ_2) from μ , we define the random set of “last cross points” C_P by

$$C_P = \{x \in P : \text{there is an infinite path of slope } s, \langle s, \xi \rangle \neq 0, \text{ in } (\tau_1, \tau_2) \\ \text{which hits } P \text{ for the last time at } x\}.$$

We analogously define the random set of “first cross points” A_P by

$$A_P = \{x \in P : \text{there is an infinite path of slope } s, \langle s, \xi \rangle \neq 0, \text{ in } (\tau_1, \tau_2) \\ \text{which hits } P \text{ for the first time at } x\}.$$

Lemma 7.4.5. *With the set up above, for μ -almost every (τ_1, τ_2) , both*

$$\lim_{n \rightarrow \infty} \frac{|C_P \cap [1, n]^3|}{n^2} \quad \text{and} \quad \lim_{n \rightarrow \infty} \frac{|A_P \cap [1, n]^3|}{n^2}$$

exist and are greater than 0.

Proof. The fact that the limits exist follows from the \mathbb{Z}^2 ergodic theorem applied along P .

The proofs are analogous, so we just present the proof for C_P . By Proposition 7.3.2, the μ -expected value of the slope along the component γ containing the origin in (τ_1, τ_2) is $s(\mu_1) - s(\mu_2)$. If γ is a double edge or finite cycle then the slope along γ is 0, so the set S of pairs of

tilings (τ_1, τ_2) such that there is an infinite path with slope in the set $\{s : \langle s, \xi \rangle \neq 0\}$ through the origin has $\mu(S) > p$ for some $p > 0$.

Since μ is ergodic with respect to the $\mathbb{Z}_{\text{even}}^3$ action, it follows that along any $\mathbb{Z}_{\text{even}}^3$ -orbit, the proportion of the orbit in S is $> p$. On the other hand, an infinite path with slope in $\{s : \langle s, \xi \rangle \neq 0\}$ only crosses P finitely many times almost surely. In particular, for any $\delta > 0$, there exists M such that

$$\mu(\ell \text{ is an infinite path passing through the origin with slope } \langle s(\ell), \xi \rangle \neq 0 \\ \text{and hits } P \text{ more than } M \text{ times}) < \delta.$$

Therefore

$$\mu\left(\lim_{n \rightarrow \infty} \frac{|C_P \cap [1, n]^3|}{n^2} > \frac{p}{M}\right) \geq 1 - \delta,$$

which completes the proof. \square

The next technical lemma is about the distribution of the distance between *hit points*. Given a plane P , let $\alpha \subset (\tau_1, \tau_2)$ be an arc of a path (finite or infinite) between two points in P , such that α is disjoint from P except its endpoints $x_\alpha, y_\alpha \in P$. We define the *distance between hits* by

$$d_P(\alpha) = \text{dist}(x_\alpha, y_\alpha)$$

where dist denotes L^1 distance on P .

Lemma 7.4.6. *Let μ be an ergodic coupling of ergodic measures μ_1, μ_2 on Ω . Let $B_n = [1, n]^3$. For any $\beta > 0$, there exists M such that for all $\theta > 0$, there exists N such that if $n \geq N$, then*

$$\mu\left(\#\{\alpha \text{ arc of path hitting } P \cap B_n : d_P(\alpha) > M\} \leq \beta n^2\right) > 1 - \theta.$$

Proof. As there is some probability distribution on the distance between hit points, by $\mathbb{Z}_{\text{even}}^3$ -invariance given $\epsilon > 0$ there exists M large enough such that for all $v \in P$,

$$\mu(x_\alpha = v, d_P(\alpha) > M) < \epsilon. \quad (49)$$

For a set of points $A \subset \mathbb{Z}^3$ let $\text{even}(A), \text{odd}(A)$ denote the subset of even, odd points respectively, and define

$$S_n^{\text{even}} = \frac{2}{n^2} \left[\#\{\alpha \subset (\tau_1, \tau_2) : x_\alpha \in \text{even}(B_n \cap P), d_P(\alpha) > M\} \right], \\ S_n^{\text{odd}} = \frac{2}{n^2} \left[\#\{\alpha \subset (\tau_1, \tau_2) : x_\alpha \in \text{odd}(B_n \cap P), d_P(\alpha) > M\} \right].$$

By the $\mathbb{Z}_{\text{even}}^2$ ergodic theorem applied along P , μ -almost everywhere S_n^{even} converges to a limit S^{even} as $n \rightarrow \infty$ (and similarly for S_n^{odd}). Further, we get that

$$\mu(x_\alpha = 0, d_P(\alpha) > M) = \int_{\Omega \times \Omega} S^{\text{even}}(\tau_1, \tau_2) d\mu(\tau_1, \tau_2) \quad (50)$$

and if v is an odd point,

$$\mu(x_\alpha = v, d_P(\alpha) > M) = \int_{\Omega \times \Omega} S^{\text{odd}}(\tau_1, \tau_2) d\mu(\tau_1, \tau_2). \quad (51)$$

Since $S_n^{\text{even}}, S_n^{\text{even}} \geq 0$, Equation (49) and Equation (50) (and analogously Equations (49), (51) for the odd case) combine to show that for n large enough,

$$\mu(S_n^{\text{even}} \leq 3\epsilon) \geq 1 - 2\epsilon \quad \text{and} \quad \mu(S_n^{\text{odd}} \leq 3\epsilon) \geq 1 - 2\epsilon.$$

Putting together the even and odd cases, for n large enough,

$$\mu\left(\#\{\alpha \text{ arc of path hitting } P \cap B_n : d_P(\alpha) > M\} \leq 3\epsilon n^2/2\right) \geq 1 - 2\epsilon.$$

Choosing ϵ appropriately given β, θ completes the proof. \square

We can now state and prove the theorem about the effect of chain swapping (with probability $p = 1/2$) on the Gibbs property.

Theorem 7.4.7. *Let ν be a Gibbs measure on $\Omega \times \Omega$ which is an ergodic coupling of ergodic measures $\nu_1 \in \mathcal{P}_e^{s_1}$ and $\nu_2 \in \mathcal{P}_e^{s_2}$ with $s_1 \neq s_2$ and $(s_1 + s_2)/2 \in \text{Int}(\mathcal{O})$. The measure ν' obtained from ν by chain swapping with probability $p = 1/2$ is not a Gibbs measure on $\Omega \times \Omega$.*

Remark 7.4.8. The condition $(s_1 + s_2)/2 \in \text{Int}(\mathcal{O})$ is necessary in the proof so that we can use the patching theorem (Theorem 6.3.5).

Proof. Let $B_n = [1, n]^3$. Let P be a coordinate plane with normal vector denoted ξ such that $P \cap B_n$ is a face of ∂B_n (denoted F) and such that $\langle s(\nu_1) - s(\nu_2), \xi \rangle \neq 0$. By Lemma 7.4.5,

$$\nu\left(\lim_{n \rightarrow \infty} \frac{|A_P \cap B_n|}{n^2} > 0\right) = 1.$$

Recall that A_P is the collection of points $x \in P$ on infinite paths $\ell \subset (\tau_1, \tau_2)$ with $\langle s(\ell), \xi \rangle \neq 0$ such that x is the *first* time that ℓ intersects P . Given a sample (τ_1, τ_2) from ν , we look at the collection of infinite paths ℓ satisfying $\langle s(\ell), \xi \rangle \neq 0$.

The part of ℓ outside B_n , $\ell \setminus B_n$, always has exactly two infinite components, a *left ray* (half-infinite path entering B_n) and a *right ray* (half-infinite path exiting B_n). We define *first entrance points* of B_n by

$$S_{\text{first}}(B_n) = \{x \in \partial B_n : \text{there is an infinite path } \ell \subset (\tau_1, \tau_2) \text{ with } \langle s(\ell), \xi \rangle \neq 0, \\ \ell \text{ enters } B_n \text{ for the first time at } x\}.$$

Note that left rays hit ∂B_n at first entrance points. Similarly define *last exit points* of B_n by

$$X_{\text{last}}(B_n) = \{x \in \partial B_n : \text{there is an infinite path } \ell \subset (\tau_1, \tau_2) \text{ with } \langle s(\ell), \xi \rangle \neq 0, \\ \ell \text{ exits } B_n \text{ for the last time at } x\}.$$

Right rays hit ∂B_n at last exit points. See Figure 30 for an illustration. We show that without loss of generality (i.e. up to translating P) there are many left rays incident to the face $F = P \cap \partial B_n$, in particular that it contains many points in $S_{\text{first}}(B_n)$. To do this, let \tilde{B}_n be B_n reflected over P and notice that

$$A_P \cap B_n \subset S_{\text{first}}(B_n) \cup S_{\text{first}}(\tilde{B}_n).$$

Therefore at least one of $A_P \cap S_{\text{first}}(B_n)$ and $A_P \cap S_{\text{first}}(\tilde{B}_n)$ has size of order n^2 . (It is possible for only one to have order n^2 points, for example if all paths in (τ_1, τ_2) are in the




Figure 30: Example of an infinite path $\ell \subset (\tau_1, \tau_2)$ hitting B_n , with first entrance, last exit, and left and right rays labeled.

same direction.) Without loss of generality (by translating and possibly changing the choice of face F), there exists $c \in (0, 1)$ such that given $\delta > 0$, for n large enough

$$|A_P \cap S_{\text{first}}(B_n)| > cn^2 \quad (52)$$

with ν -probability $1 - \delta$.

Given $x \in A_P \cap S_{\text{first}}(B_n)$, there exists a unique infinite path $\ell \subset (\tau_1, \tau_2)$ with slope denoted $s(\ell)$ containing x . Since $x \in A_P$ this path will have $\langle s(\ell), \xi \rangle \neq 0$, so ℓ hits P finitely many times almost surely. We define the function $D_P(\ell)$ to be the distance along P from $\ell \cap A_P$ to $\ell \cap C_P$ (note that this is different from $d_P(\cdot)$ defined in Lemma 7.4.6).

Without loss of generality assume that the origin is contained in P . Let ℓ_0 be the path through the origin in (τ_1, τ_2) . Then for any $\theta > 0$ there exists M such that

$$\nu(D_P(\ell_0) > M \mid \langle s(\ell), \xi \rangle \neq 0) < \theta. \quad (53)$$

By $\mathbb{Z}_{\text{even}}^3$ -invariance, this holds for any ℓ through an even point on P with $\langle s(\ell), \xi \rangle \neq 0$. An analogous statement to Equation (53) holds if we look at an odd point $v \in P$, and $\mathbb{Z}_{\text{even}}^3$ -invariance again implies that it holds for any ℓ through an odd point on P with $\langle s(\ell), \xi \rangle \neq 0$. Putting these together, we have that for n large

$$\nu(\#\{\ell : \ell \cap C_P \notin B_n, \ell \cap A_P \in B_n\} \leq 4Mn + \theta n^2) > 1 - \theta.$$

Take $M = \epsilon n$, with $\epsilon > 0$ small to be specified below. Then for n large this becomes

$$\nu(\#\{\ell : \ell \cap C_P \notin B_n, \ell \cap A_P \in B_n\} \leq (4\epsilon + \theta)n^2) > 1 - \theta.$$

As any infinite path $\ell \subset (\tau_1, \tau_2)$ has well-defined slope, if $\langle s(\ell), \xi \rangle \neq 0$ then ℓ must be on opposite sides of P before A_P and after C_P . Hence by the above and Equation (52), for n large enough,

$$\nu(\#\{\ell : \ell \cap S_{\text{first}}(B_n) \in P, \ell \cap X_{\text{last}}(B_n) \notin P\} > (c - 4\epsilon - \theta)n^2) > 1 - \theta - \delta. \quad (54)$$

Therefore with ν -probability $1 - \theta - \delta$, at least $c'n^2 = (c - 4\epsilon - \theta)n^2$ infinite paths ℓ entering at $x \in A_P \cap B_n \subset F$ exit B_n at $y \notin F$.

On the other hand, since $s_1 \neq s_2$, we can apply chain swapping with $p = 1/2$ to get a new measure ν' distinct from ν . By Proposition 7.4.2, the marginals ν'_1, ν'_2 of ν' are ergodic. By Proposition 7.4.4, they satisfy

$$s(\nu'_1) = s(\nu'_2) = \frac{s_1 + s_2}{2}.$$

Together this means that ν'_1, ν'_2 satisfy the conditions of the patching theorem (Theorem 6.3.5). Fixing $\epsilon \in (0, 1)$, let A_n be the cubic annulus between B_n and $(1 - \epsilon)B_n$. By Theorem 6.3.5 applied to ν'_1, ν'_2 on A_n , for (τ'_1, τ'_2) sampled from ν' , for n large enough we can with ν' -probability $1 - \epsilon$ find a tiling τ such that

- $\tau|_{(1-\epsilon)B_n} = \tau'_2$
- $\tau|_{\mathbb{Z}^3 \setminus B_n} = \tau'_1$.

Let $Z_n \subset A_P \cap S_{\text{first}}(B_n)$ be the subset of points x such that the infinite path ℓ through x in (τ_1, τ_2) satisfies:

- ℓ has $C_P \cap \ell \in B_n$ (so that ℓ exits B_n through $\partial B_n \setminus P$);
- ℓ did not have its orientation reversed by the chain swapping (in other words, $\ell \subset (\tau'_1, \tau'_2) \cap (\tau_1, \tau_2)$).

By Equation (54) and since chain swapping reverses the orientation of each infinite path with independent probability $1/2$, given $\delta' > 0$, for n large enough, setting $c' = c - 4\epsilon - \theta$ we have

$$\nu'(|Z_n| > c'n^2/2) > \frac{1}{2} - \delta' > 0. \quad (55)$$

Conditional on the double dimer configuration (τ'_1, τ'_2) on $\mathbb{Z}^3 \setminus B_n$, if ν' is a Gibbs measure then it must assign the same probability to (τ'_1, τ'_2) and (τ, τ'_2) . However since τ, τ'_2 agree on $(1 - \epsilon)B_n$, there are no infinite paths in (τ, τ'_2) through $(1 - \epsilon)B_n$.

Let $S'_{\text{first}}(B_n)$ and $X'_{\text{last}}(B_n)$ denote the first entrance and last exit points in (τ'_1, τ'_2) . We note that

$$S'_{\text{first}}(B_n) \cup X'_{\text{last}}(B_n) = S_{\text{first}}(B_n) \cup X_{\text{last}}(B_n)$$

because on infinite paths where the orientation was swapped, the first entrance and last exit points are swapped.

On the other hand, since (τ'_1, τ'_2) and (τ, τ'_2) agree on $\mathbb{Z}^3 \setminus B_n$, they have the same first entrance and last exit points and the same left and right rays. If $x \in S'_{\text{first}}(B_n)$ and $y \in X'_{\text{last}}(B_n)$, we denote the left and right rays incident to them by $\ell_-(x)$ and $\ell_+(y)$ respectively. The tiling $(\tau, \tau'_2)|_{B_n}$ pairs up all the left rays with right rays in a new way to make full infinite paths.

However recall that an infinite path in a double dimer configuration sampled from ν' has well-defined slope almost surely. We show that ν' is *not* Gibbs by showing that it is *not* possible to pair order n^2 of the left rays entering at $x \in Z_n$ with right rays of the same slope. For $x \in Z_n$, let $\gamma(x) \subset (\tau, \tau'_2)$ denote the path that connects $\ell_-(x)$ to an exit point y . Then the infinite path in (τ, τ'_2) through x is

$$\ell_-(x) \cup \gamma(x) \cup \ell_+(y).$$




Figure 31: In all three pictures, the transparent cube is B_n and smaller orange cube inside it is $(1 - \epsilon)B_n$. The front left face is F . In (τ, τ'_2) , all tiles in $(1 - \epsilon)B_n$ are double edges, so infinite paths can't enter the orange cube. The region T , corresponding to Case 1, is the union of the three blue regions.

The remainder of the proof is casework to show there only a small number of these infinite paths can have well-defined slope. Recall that $F = B_n \cap P$ and let F° denote the points in F which are distance $\geq \epsilon n$ from ∂F .

1. **Bounded by area:** We define the *thin region* T , which is a union of three things: i) $F \setminus F^\circ$, ii) $F \setminus F^\circ$ translated ϵn inward, iii) the part of ∂B_n between i) and ii). See Figure 31 for an illustration. Since

$$\text{area}(T) \leq 12\epsilon n^2,$$

the number of infinite paths in (τ, τ'_2) which intersect T is bounded by $12\epsilon n^2$.

2. **Bounded by number of possible connecting paths:** choose $x \in Z_n$, and suppose that $\ell_-(x) \cup \gamma(x) \cup \ell_+(y)$ does not intersect T . To have well-defined slope, ℓ must still cross P some time after x , and to avoid T it must at some point cross P in $P \setminus F$. See Figure 32.

Therefore the rest of the path $\gamma(x) \cup \ell_+(y)$ must use part of at least one finite cycle or infinite path in

$$(\tau, \tau'_2) |_{\mathbb{Z}^3 \setminus B_n} = (\tau'_1, \tau'_2) |_{\mathbb{Z}^3 \setminus B_n}$$

to connect a point in F° to $P \setminus F$. This path will be an arc on P in the P -half-space on the opposite side of P from B_n . Chain swapping only changes the directions of paths, so the collection of arcs and their lengths are the same in (τ_1, τ_2) and (τ'_1, τ'_2) . Since the arcs are outside B_n , they are also the same in (τ, τ'_2) . Thus by Lemma 7.4.6 applied to $M = \epsilon n$, we can find $\beta, \theta > 0$ small enough so that for n large enough,

$$\nu' \left(\#\{\alpha \text{ arc of path hitting } F : d_P(\alpha) > \epsilon n\} < \beta n^2 \right) > 1 - \theta,$$

where recall that $d_P(\alpha)$ is the distance along P between the two intersection points of the arc α with P . Therefore with ν' -probability $1 - \theta$, the number of $x \in Z_n$ such that the path

$$\ell_-(x) \cup \gamma(x) \cup \ell_+(y)$$

is disjoint from T but crosses $P \setminus F^\circ$ is at most βn^2 .




Figure 32: Two examples corresponding to Case 2. In this case the infinite path does not intersect T , so this can happen either if the final exit point $y \in \partial B_n \setminus F$ (left) or if the final exit point $y \in F$, but the right ray $\ell_+(y)$ crosses P again outside B_n (right).




Figure 33: Corresponding to Case 3, if $\gamma(x) \cup \ell_+(y)$ never crosses $P \setminus F$, then the resulting infinite path cannot have well-defined slope.

3. **Remaining paths forced to have no well-defined slope:** if $x \in Z_n$ is not in Case 1 or Case 2, then the path $\ell := \ell_-(x) \cup \gamma(x) \cup \ell_+(y)$ does not intersect T and does not cross $P \setminus F^\circ$ at any time after going through x . This implies that $\ell_-(x)$ and $\ell_+(y)$ are contained in the same P half-space, in which case ℓ cannot have well-defined slope. See Figure 33.

In summary, with probability $1 - \theta$, there are at most

$$(12\epsilon + \beta)n^2$$

points $x \in Z_n$ such that we can connect $\ell_-(x)$ to have well-defined slope. However by Equation (55), with positive ν -probability $|Z_n| > c'n^2/2$. We can take ϵ, β, θ to be arbitrarily small compared to c , and thus since ν' is $\mathbb{Z}_{\text{even}}^3$ -invariant (meaning infinite paths must have well-defined slope a.s.), $\nu'(\cdot | (\tau'_1, \tau'_2) |_{\mathbb{Z}^3 \setminus B_n})$ cannot assign the same probability to (τ'_1, τ'_2) and (τ, τ'_2) . Therefore ν' is not a Gibbs measure. \square

Theorem 7.4.7 and Corollary 7.1.5 combine to give the following corollary.

Corollary 7.4.9. *If μ is a Gibbs measure on $\Omega \times \Omega$ which is an ergodic coupling of $\mu_1 \in \mathcal{P}_e^{s_1}$*

and $\mu_2 \in \mathcal{P}_e^{s_2}$ for $s_1 \neq s_2$ and $\frac{s_1+s_2}{2} \in \text{Int}(\mathcal{O})$, then the measure μ' obtained by chain swapping with probability $p = 1/2$ does not maximize entropy in $\mathcal{P}^{\frac{s_1+s_2}{2}, \frac{s_1+s_2}{2}}$.

7.5 Strict concavity of ent and existence of EGMs of every mean current

With the chain swapping machinery developed in Section 7.4, we can now prove one of the main results of this section, namely that that ent is *strictly* concave on $\mathcal{O} \setminus \mathcal{E}$ (recall that \mathcal{E} denotes the edges of \mathcal{O}). We already showed that ent is concave on \mathcal{O} in Lemma 7.2.1. We also already showed in Section 4 that ent restricted to the interior of any face of $\partial\mathcal{O}$ is strictly concave (Corollary 4.2.8) by relating ent restricted to a face of $\partial\mathcal{O}$ to ent_{loz} , the slope entropy function for two-dimensional lozenge tilings (Theorem 4.2.7).

Theorem 7.5.1. *The entropy function ent is strictly concave on $\mathcal{O} \setminus \mathcal{E}$.*

Proof. By Lemmas 7.2.1 and 7.2.3, ent is concave and continuous on \mathcal{O} . To show strict concavity on $\mathcal{O} \setminus \mathcal{E}$, it suffices to show that if $s_1, s_2 \in \mathcal{O}$ and $(s_1 + s_2)/2 \in \mathcal{O} \setminus \mathcal{E}$, then

$$\text{ent}((s_1 + s_2)/2) > (\text{ent}(s_1) + \text{ent}(s_2))/2.$$

If $(s_1 + s_2)/2$ is contained in the interior of a face $\mathcal{F} \subset \partial\mathcal{O}$, then we are done by Corollary 4.2.8. The remaining case is that $(s_1 + s_2)/2 \in \text{Int}(\mathcal{O})$. In this case let μ_1 and μ_2 be entropy maximizers in \mathcal{P}^{s_1} and \mathcal{P}^{s_2} respectively (these exist by Lemma 7.1.1) and let μ be the independent coupling of μ_1 and μ_2 . Then

$$\text{ent}(s_1) + \text{ent}(s_2) = h(\mu).$$

Consider the ergodic decomposition

$$\mu = \int \nu \, dw_\mu(\nu)$$

where w_μ is a probability measure on the space of ergodic couplings of ergodic Gibbs measures (see Proposition 2.2.1, which says that the ergodic components of a Gibbs measure are Gibbs a.s.). Let μ' be the measure obtained by applying chain swapping with probability $p = 1/2$ to μ . By Proposition 7.4.2,

$$\mu' = \int \nu' \, dw_\mu(\nu)$$

where ν' is obtained from ν by chain swapping also with $p = 1/2$. By Proposition 7.4.3, $h(\nu) = h(\nu')$. Since $h(\cdot)$ is an affine function, we get that

$$h(\mu) = \int h(\nu) \, dw_\mu(\nu) = \int h(\nu') \, dw_\mu(\nu) = h(\mu').$$

By Proposition 7.4.4, $s(\pi_1(\nu')) = s(\pi_2(\nu')) = (s(\nu_1) + s(\nu_2))/2$. Since $s(\cdot)$ is an affine function,

$$s(\pi_1(\mu')) = s(\pi_2(\mu')) = (s_1 + s_2)/2.$$

Let $(\mathfrak{s}_1, \mathfrak{s}_2)$ denote the random pair of mean currents for a double dimer configuration sampled from μ . To complete the proof, we proceed with cases based on w_μ . Consider the sets

$$A = \left\{ \nu : \frac{s(\pi_1(\nu)) + s(\pi_2(\nu))}{2} \in \text{Int}(\mathcal{O}) \right\}, \quad B = \left\{ \nu : s(\pi_1(\nu)) \neq s(\pi_2(\nu)) \right\}.$$

If $\nu \in A \cap B$ is an ergodic coupling of ergodic measures, then ν satisfies the conditions of Theorem 7.4.7. Since $\mathbb{E}[\mathfrak{s}_1 - \mathfrak{s}_2] = s_1 - s_2 \neq 0$, $w_\mu(B) > 0$. If $w_\mu(A) > 0$, then since μ is an independent coupling, we can argue in a few cases that $w_\mu(A \cap B) > 0$:

- If $\mathfrak{s}_1, \mathfrak{s}_2$ are both atomic, then $w_\mu(A \cap B) > 0$. For the next cases we assume without loss of generality that \mathfrak{s}_2 is not atomic.
- If $\{\mathfrak{s}_1 \in \text{Int}(\mathcal{O})\}$ has positive probability, then given any value of \mathfrak{s}_1 in $\text{Int}(\mathcal{O})$, \mathfrak{s}_2 has positive probability to be different from it. Since $\mathfrak{s}_1 \in \text{Int}(\mathcal{O})$, the average is in $\text{Int}(\mathcal{O})$.
- If \mathfrak{s}_1 has positive probability to be contained in $\partial\mathcal{O}$, then given any value of \mathfrak{s}_1 in $\partial\mathcal{O}$, $w_\mu(A) > 0$ implies that \mathfrak{s}_2 has positive probability to not be contained in the same face as \mathfrak{s}_1 (since on A , their average must be in $\text{Int}(\mathcal{O})$). On the other hand if \mathfrak{s}_2 is not contained in the same face of $\partial\mathcal{O}$ as \mathfrak{s}_1 , then it must be different from \mathfrak{s}_1 .

Applying Theorem 7.4.7 shows that ν' is not a Gibbs measure for $\nu \in A \cap B$. Since the ergodic components of Gibbs measures are Gibbs a.s., if $w_\mu(A \cap B) > 0$ then μ' is not a Gibbs measure. By Corollary 7.1.5, μ' is not an entropy maximizer in $\mathcal{P}^{\frac{s_1+s_2}{2}, \frac{s_1+s_2}{2}}$, and hence

$$\text{ent}(s_1) + \text{ent}(s_2) = h(\mu) = h(\mu') < 2\text{ent}((s_1 + s_2)/2).$$

This completes the proof if $w_\mu(A) > 0$.

However it can happen that $w_\mu(A) = 0$ (for example, if \mathfrak{s}_1 is supported at a corner vertex $v \in \partial\mathcal{O}$, and \mathfrak{s}_2 is supported on a square on $\partial\mathcal{O}$ around v). Since μ is an independent coupling, $w_\mu(A) = 0$ implies that $\mathfrak{s}_1, \mathfrak{s}_2$ are supported in $\partial\mathcal{O}$. There are two remaining cases.

First suppose there is a face $\mathcal{F} \subset \partial\mathcal{O}$ such that

$$C = \left\{ \nu : \frac{s(\pi_1(\nu)) + s(\pi_2(\nu))}{2} \in \text{Int}(\mathcal{F}) \right\}$$

has $w_\mu(C) > 0$. Since μ is an independent coupling, $w_\mu(C \cap B) > 0$ by arguments analogous to those above for A, B . Let μ' be obtained from μ by chain swapping. By Proposition 7.4.2,

$$\mu' = \int \nu' dw_\mu(\nu)$$

where ν' is obtained from ν by chain swapping. Let ν_1, ν_2 denote the marginals of ν and let ν'_1, ν'_2 denote the marginals of ν' . If $\nu \in B$, then $s(\nu_1) \neq s(\nu_2)$ and ν' is distinct from ν . By Proposition 7.4.3 and Proposition 7.4.4,

$$h(\nu') = h(\nu), \quad s(\nu'_1) = s(\nu'_2) = \frac{s(\nu_1) + s(\nu_2)}{2}.$$

By Theorem 4.2.7, $\text{ent}|_{\mathcal{F}} = \text{ent}_{\text{loz}}$. Since ent_{loz} is strictly concave on $\text{Int}(\mathcal{F})$, we have that for each $\nu \in C \cap B$,

$$h(\nu') = h(\nu'_1) + h(\nu'_2) < 2\text{ent}((s(\nu_1) + s(\nu_2))/2)$$

Since $w_\mu(C \cap B) > 0$, Lemma 7.2.1 and the affine property of h implies that

$$\text{ent}(s_1) + \text{ent}(s_2) = h(\mu) = h(\mu') = \int h(\nu) dw_\mu(\nu) < 2\text{ent}((s_1 + s_2)/2).$$

This completes the proof in the case that $w_\mu(C) > 0$.

Finally if $w_\mu(A) = 0$ and $w_\mu(C) = 0$ for all faces of $\partial\mathcal{O}$, then $\mathfrak{s}_1, \mathfrak{s}_2$ must be supported in \mathcal{E} (for example, \mathfrak{s}_1 could be supported at one vertex $v \in \partial\mathcal{O}$, and \mathfrak{s}_2 could be supported on the four edges of $\partial\mathcal{O}$ incident to v). Since $\text{ent}|_{\mathcal{E}} \equiv 0$, this implies that $h(\mu_1) + h(\mu_2) = h(\mu) = 0$ and hence that $h(\mu_1) = h(\mu_2) = 0$.

However by Lemma 7.2.1 and Theorem 4.2.7, $\text{ent}|_{\mathcal{O} \setminus \mathcal{E}} > 0$. Therefore if $s_1, s_2 \in \mathcal{O} \setminus \mathcal{E}$, then μ_1, μ_2 cannot be entropy maximizers in $\mathcal{P}^{s_1}, \mathcal{P}^{s_2}$. This completes the proof. \square

With this we can strengthen Theorem 7.1.2.

Theorem 7.5.2. *For every $s \in \text{Int}(\mathcal{O})$, a measure $\mu \in \mathcal{P}^s$ satisfies $h(\mu) = \text{ent}(s)$ if and only if μ is a convex combination of ergodic Gibbs measures of mean current s . In particular, if $\nu \in \mathcal{P}^s$ is an ergodic Gibbs measure, then $h(\nu) = \text{ent}(s)$.*

Remark 7.5.3. In contrast, an EGM of mean current $s \in \partial\mathcal{O}$ can have any specific entropy between 0 and $\text{ent}(s)$, see Proposition 4.2.10. (Note however that for $s \in \mathcal{E}$, $\text{ent}(s) = 0$.) All EGMs of mean current $s \in \text{Int}(\mathcal{O})$ have the same specific entropy by Corollary 6.8.1.

Proof. Suppose $\mu \in \mathcal{P}^s$ maximizes entropy (μ exists by Lemma 7.1.1). By Theorem 7.1.2, μ is a Gibbs measure. Consider its ergodic decomposition

$$\mu = \int \nu \, dw_\mu(\nu).$$

Since ergodic components of Gibbs measures are Gibbs a.s., w_μ is a probability measure on ergodic Gibbs measures. Since $h(\cdot)$ is an affine function, it follows that

$$\text{ent}(s) = h(\mu) = \int h(\nu) \, dw_\mu(\nu) \leq \int \text{ent}(s(\nu)) \, dw_\mu(\nu).$$

Since $s \in \text{Int}(\mathcal{O})$, by Theorem 7.5.1 if $s(\nu)$ is not constant then the middle inequality below is strict:

$$\text{ent}(s) \leq \int \text{ent}(s(\nu)) \, dw_\mu(\nu) < \text{ent}\left(\int s(\nu) \, dw_\mu(\nu)\right) = \text{ent}(s).$$

Therefore all ergodic components ν of μ must have $s(\nu) = s$, i.e., the support of w_μ is contained in the set of ergodic Gibbs measures of mean current s .

By Corollary 6.8.1, if ν_1, ν_2 are EGMs of the same mean current $s \in \text{Int}(\mathcal{O})$, then $h(\nu_1) = h(\nu_2)$. Therefore if $\nu \in \mathcal{P}^s$ is an ergodic Gibbs measure, $\text{ent}(s) = h(\nu)$. □

From Theorem 7.5.2 and Lemma 7.1.1 for interior mean currents and the results of Section 4 for boundary ones, there exist ergodic Gibbs measures of all mean currents.

Corollary 7.5.4. *For all $s \in \mathcal{O}$, there exists an ergodic Gibbs measure of mean current s .*

Remark 7.5.5. In two dimensions, there exists a *unique* ergodic Gibbs measure of every interior slope. Uniqueness of EGMs for interior mean currents is open problem, see Problem 9.0.3 and the related Problem 9.0.4.

7.6 Properties of Ent

Recall that $AF(R)$ denotes the space of asymptotic flows on R , and $AF(R, b)$ denotes the asymptotic flows on R with boundary value b . Both are equipped with the Wasserstein metric d_W (see Section 5). Here we use the properties of the mean-current entropy function ent to prove things about Ent , the entropy functional on asymptotic flows given by

$$\text{Ent}(f) = \frac{1}{\text{Vol}(R)} \int_R \text{ent}(f(x)) \, dx.$$

As Corollaries of Lemma 7.2.1, Theorem 7.5.1 and Lemma 7.2.3 respectively we get that

Corollary 7.6.1. *The entropy functional Ent is concave on $AF(R)$. Further, Ent is strictly concave when restricted to the space of asymptotic flows which are valued in $\mathcal{O} \setminus \mathcal{E}$.*

Corollary 7.6.2. *If $f_n \rightarrow f$ almost everywhere in R , then $\text{Ent}(f) = \lim_{n \rightarrow \infty} \text{Ent}(f_n)$.*

From this, we show

Proposition 7.6.3. *The functional $\text{Ent} : AF(R) \rightarrow [0, \infty)$ is upper semicontinuous in the Wasserstein topology induced by d_W .*

Proof. Let $(f_n)_{n \geq 1}$ be a sequence of flows in $AF(R)$ such that $d_W(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$ for some $f \in AF(R)$. For any $g \in AF(R)$, we can define its approximation g_ϵ given by

$$g_\epsilon(x) := \frac{1}{\text{Vol}B_\epsilon(x)} \int_{B_\epsilon(x)} g(y) \, dy.$$

Here we say that $g(y) = 0$ if $y \notin R$. While g_ϵ is not an asymptotic flow because it is not divergence-free, it is still valued in \mathcal{O} and thus $\text{Ent}(g_\epsilon) := \frac{1}{\text{Vol}(R)} \int_R \text{ent}(g_\epsilon(x)) \, dx$ still makes sense. By the Lebesgue differentiation theorem, g_ϵ converges to g almost everywhere as $\epsilon \rightarrow 0$. By Corollary 7.6.2,

$$\lim_{\epsilon \rightarrow 0} \text{Ent}(g_\epsilon) = \text{Ent}(g).$$

By Lemma 7.2.1, for any $x \in R$,

$$\text{ent}(g_\epsilon(x)) = \text{ent}\left(\frac{3}{4\pi\epsilon^3} \int_{B_\epsilon(x)} g(y) \, dy\right) \geq \frac{3}{4\pi\epsilon^3} \int_{B_\epsilon(x)} \text{ent}(g(y)) \, dy.$$

Therefore there is a constant C (proportional to $\text{Area}(\partial R)/\text{Vol}(R)$ and independent of ϵ) such that

$$\text{Ent}(g_\epsilon) + C\epsilon \geq \text{Ent}(g).$$

Since $d_W(f_n, f) \rightarrow 0$ as $n \rightarrow \infty$, by Corollary 5.1.3, $f_{n,\epsilon}$ converges pointwise to f_ϵ . By Corollary 7.6.2,

$$\limsup_{n \rightarrow \infty} \text{Ent}(f_n) \leq \limsup_{n \rightarrow \infty} \text{Ent}(f_{n,\epsilon}) + C\epsilon = \text{Ent}(f_\epsilon) + C\epsilon.$$

Taking ϵ to zero, we get that

$$\limsup_{n \rightarrow \infty} \text{Ent}(f_n) \leq \text{Ent}(f),$$

hence Ent is upper semicontinuous. □

Remark 7.6.4. It is not difficult to see that Ent is not continuous. Indeed consider the flows $f_n \in AF([0, 1]^3)$ given by

$$f_n(x_1, x_2, x_3) = \begin{cases} \eta_2 & \text{if } x_1 \in (\frac{2k}{2n}, \frac{2k+1}{2n}) \text{ for some } 0 \leq k \leq n-1 \\ -\eta_2 & \text{if } x_1 \in (\frac{2k+1}{2n}, \frac{2k+2}{2n}) \text{ for some } 0 \leq k \leq n-1. \end{cases}$$

Then f_n converges to the constant zero vector field but $\text{Ent}(f_n) = 0$ while $\text{Ent}(0) > 0$.

Our main goal is to show that there exists a unique Ent maximizer in $AF(R, b)$ under some mild conditions on the pair (R, b) . Standard analytic arguments are enough to show existence and a weak form of uniqueness. Let $\epsilon_1, \dots, \epsilon_8$ denote the eight closed edges of \mathcal{O} which make up \mathcal{E} .

Proposition 7.6.5. *There exists $f \in AF(R, b)$ such that $\text{Ent}(f) = \sup_{g \in AF(R, b)} \text{Ent}(g)$. Further, given $f_1, f_2 \in AF(R, b)$, define*

$$A = \{x \in R : f_1(x) \neq f_2(x)\}, \quad B = \bigcup_{i=1}^8 \{x \in R : f_1(x), f_2(x) \in \mathfrak{e}_i\}.$$

If f_1, f_2 are both Ent maximizers, then $A \subseteq B$.

Remark 7.6.6. The problem is that ent is only strictly concave on $\mathcal{O} \setminus \mathcal{E}$, not all of \mathcal{O} . The same problem arises in two dimensions, and is addressed in [Gor21] and [DSS10].

Proof. Since $(AF(R, b), d_W)$ is compact (Theorem 5.3.4) and Ent is upper semicontinuous (Proposition 7.6.3), the existence of the maximizer follows.

To prove weak uniqueness, recall that $\text{ent}(s) = 0$ if and only if $s \in \mathcal{E}$. If f_1, f_2 are distinct maximizers then A has positive measure. If $A \cap (R \setminus B)$ has positive measure, then by strict convexity of Ent on flows valued in $\mathcal{O} \setminus \mathcal{E}$ (Corollary 7.6.1),

$$\text{Ent}\left(\frac{f_1 + f_2}{2}\right) > \text{Ent}(f_1) + \text{Ent}(f_2),$$

which would contradict the claim that f_1, f_2 are maximizers. Therefore $A \subseteq B$. □

We adapt an argument of V. Gorin in [Gor21, Proposition 7.10] to prove uniqueness under the mild condition that the pair (R, b) is *semi-flexible* as defined in Definition 7.6.8 below. We call this *semi-flexible* since it is a weaker condition than *flexible*, which will be defined at the beginning of Section 8.

Definition 7.6.7. Fix a boundary asymptotic flow b on R . A point $x \in R$ with boundary condition b is *frozen* if for all open sets $U \ni x$ and all entropy maximizers $f \in AF(R, b)$, there are points $y \in U$ such that $f(y) \in \mathcal{E}$. A point $x \in R$ with boundary condition b is *always frozen* if for all open sets $U \ni x$ and all $g \in AF(R, b)$, there are points $y \in U$ such that $f(y) \in \mathcal{E}$.

Definition 7.6.8. The pair (R, b) is *semi-flexible* if there are no always frozen points in $\text{Int}(R)$. I.e., (R, b) is semi-flexible if for all $x \in \text{Int}(R)$, there exists an extension $g \in AF(R, b)$ and an open set $U \ni x$ such that $g(U) \subset \mathcal{O} \setminus \mathcal{E}$. If (R, b) is not semi-flexible, we say (R, b) is *rigid*.

Remark 7.6.9. The weak uniqueness statement in Proposition 7.6.5 can be rephrased as saying that entropy maximizers are unique on the complement of the frozen points. In particular the task that remains is to show that a region (i.e. the set of frozen points) cannot both be frozen and have multiple tilings.

Theorem 7.6.10. *If (R, b) is semi-flexible, then there is a unique Ent maximizer in $AF(R, b)$.*

Remark 7.6.11. We do not know of an example of a three-dimensional region $R \subset \mathbb{R}^3$ with boundary value b such that (R, b) is rigid, so we do not know of an example in our context where the maximizer is not unique. However see Problem 9.0.7, which includes a two-dimensional, non-planar example where the maximizer is not unique.

To prove Theorem 7.6.10, we show that an equivalent definition of (R, b) semi-flexible is that b has an extension f_0 valued in $\mathcal{O} \setminus \mathcal{E}$ on $\text{Int}(R)$ (Lemma 7.6.14). After that, the key step is to show that if a maximizer takes values in \mathcal{E} , we can perturb it by f_0 to get a flow which does not take edge values and has more entropy (Lemma 7.6.15). In particular we have the corollary that even if uniqueness fails for (R, b) , it holds if b is replaced by (say) $.999b$.

Corollary 7.6.12. *Given any boundary asymptotic flow b on R and any $\delta \in (0, 1)$ there is a unique entropy maximizer in $AF(R, \delta b)$.*

Remark 7.6.13. It is also not hard to see directly that $(R, \delta b)$ is semi-flexible, and in fact flexible, see Definition 8.2.1 and Remark 8.2.3.

Lemma 7.6.14. *The pair (R, b) is semi-flexible if and only if there exists $f_0 \in AF(R, b)$ such that f_0 is valued in $\mathcal{O} \setminus \mathcal{E}$.*

Proof. The reverse implication is clear, since for any $x \in \text{Int}(R)$, taking U small enough so that $U \subset \text{Int}(R)$, f_0 is an extension such that $f_0(y) \notin \mathcal{E}$ for all $y \in U$.

If (R, b) is semi-flexible, for all $x \in \text{Int}(R)$ there exists an open set $U_x \ni x$ and $f_x \in AF(R, b)$ such that $f_x(y) \in \mathcal{O} \setminus \mathcal{E}$ for all $y \in U_x$. If $U'_x \subset U_x$ is a smaller open set, then clearly the same property holds for U'_x .

Let $\{V_i\}_{i \in \mathbb{N}}$ be the collection of open balls centered at rational points in R with rational radii. For any pair (x, U_x) we can find V_i such that $x \in V_i$ and $V_i \subset U_x$. Therefore for each $i \in \mathbb{N}$, there exists $g_i \in AF(R, b)$ such that g_i is valued in $\mathcal{O} \setminus \mathcal{E}$ on V_i . Hence the flow

$$f_0 := \sum_{i=1}^{\infty} \frac{1}{2^i} g_i$$

is valued in $\mathcal{O} \setminus \mathcal{E}$ everywhere in $\text{Int}(R)$ as desired. \square

We follow the same strategy as in [Gor21, Proposition 7.10] to prove Theorem 7.6.10 using the nowhere-edge-valued extension f_0 . The key step is:

Lemma 7.6.15. *Suppose that (R, b) is semi-flexible, and let $\mathcal{V} \subset \mathcal{E}$ denote the vertices of $\partial\mathcal{O}$. If $f \in AF(R, b)$ maximizes Ent , then up to a set of measure zero f does not take values in $\mathcal{E} \setminus \mathcal{V}$.*

Proof. Suppose for contradiction that f is an Ent maximizer in $AF(R, b)$ which takes values in $\mathcal{E} \setminus \mathcal{V}$ on a set A of positive measure, and that f_0 is an extension of the form guaranteed by Lemma 7.6.14. We will contradict the claim that f is a maximizer by showing that perturbing f by f_0 increases Ent .

By Theorem 4.2.7, if $s = (s_1, s_2, s_3)$ is contained in a face of $\partial\mathcal{O}$ then $\text{ent}(s)$ is equal to the entropy function for two dimensional lozenge tilings, namely

$$\text{ent}(s) = \text{ent}_{\text{loz}}(|s_1|, |s_2|, |s_3|) = \frac{1}{\pi} \left(L(\pi|s_1|) + L(\pi|s_2|) + L(\pi|s_3|) \right),$$

where $L(\theta) = \int_0^\theta \log(2 \sin t) dt$ ([CKP01], see Section 4). As in the proof in two dimensions [Gor21, Proposition 7.10], note from this formula that if $s \in \mathcal{E} \setminus \mathcal{V}$ and t is contained in a face of $\partial\mathcal{O}$ adjacent to the edge containing s , then for $\epsilon > 0$ small enough

$$\text{ent}(\epsilon t + (1 - \epsilon)s) > c\epsilon \log(1/\epsilon)$$

for some constant $c > 0$ depending on s, t . More generally, $t \in \text{Int}(\mathcal{O})$ can be written as a weighted average of the six brickwork patterns. Simplifying a bit, this means that t can be written as a weighted average of t_1, t_2 in the faces adjacent to the edge containing s (this takes into account four brickwork patterns), and t_3 in the edge diagonally opposite the edge containing s (this takes into account the remaining two).

$$t = \alpha t_1 + \beta t_2 + \gamma t_3, \quad \alpha + \beta + \gamma = 1.$$

By strict concavity of ent on $\mathcal{O} \setminus \mathcal{E}$ (Theorem 7.5.1),

$$\text{ent}(\epsilon t + (1 - \epsilon)s) > \alpha \text{ent}(\epsilon t_1 + (1 - \epsilon)s) + \beta \text{ent}(\epsilon t_2 + (1 - \epsilon)s) + \gamma \text{ent}(\epsilon t_3 + (1 - \epsilon)s). \quad (56)$$

We can use the two-dimensional result directly to bound the first two terms from below. For the third term, we note that $\epsilon t_3 + (1 - \epsilon)s \in \text{Int}(\mathcal{O})$, and for ϵ small enough this whole quantity can be written as an average of mean currents on the faces adjacent to the edge containing s . Using strict concavity of ent on $\mathcal{O} \setminus \mathcal{E}$ we can again apply the lower bound from the two-dimensional result. In summary, for $\epsilon > 0$ small enough, there is a constant $c > 0$ depending on s, t so that

$$\text{ent}(\epsilon t + (1 - \epsilon)s) > c\epsilon \log(1/\epsilon). \quad (57)$$

We now consider the perturbation

$$(1 - \epsilon)f + \epsilon f_0 \in AF(R, b).$$

Let $M = \sup_{s \in \mathcal{O}} \text{ent}(s)$ (this is finite because ent is continuous). For all $x \in R \setminus A$, since ent is concave on all of \mathcal{O} (Lemma 7.2.1) and non-negative we have

$$\text{ent}((1 - \epsilon)f(x) + \epsilon f_0(x)) - \text{ent}(f(x)) \geq \epsilon(\text{ent}(f_0(x)) - \text{ent}(f(x))) \geq -M\epsilon.$$

Therefore

$$\int_{R \setminus A} \text{ent}((1 - \epsilon)f(x) + \epsilon f_0(x)) dx - \int_{R \setminus A} \text{ent}(f(x)) dx \geq -M\epsilon \text{Vol}(R \setminus A).$$

On the other hand by Equation (57), for ϵ small enough there exists $A' \subset A$ of positive measure and a fixed constant $c > 0$ such that for all $x \in A'$,

$$\text{ent}((1 - \epsilon)f(x) + \epsilon f_0(x)) - \text{ent}(f(x)) = \text{ent}((1 - \epsilon)f(x) + \epsilon f_0(x)) > c\epsilon \log(1/\epsilon).$$

Therefore

$$\text{Ent}((1 - \epsilon)f + \epsilon f_0) - \text{Ent}(f) \geq \frac{-M\epsilon \text{Vol}(R \setminus A) + \text{Vol}(A')c\epsilon \log(1/\epsilon)}{\text{Vol}(R)}.$$

For $\epsilon > 0$ small enough this implies $\text{Ent}((1 - \epsilon)f + \epsilon f_0) > \text{Ent}(f)$ and contradicts the claim that f is an entropy maximizer. \square

Proof of Theorem 7.6.10. Suppose that f_1, f_2 are maximizers of Ent in $AF(R, b)$. By Lemma 7.6.15, they cannot take values in $\mathcal{E} \setminus \mathcal{V}$. By Proposition 7.6.5, they can only differ on frozen points, so

$$\{x \in R : f_1(x) \neq f_2(x)\} \subseteq \{x \in R : f_1(x), f_2(x) \in \mathcal{V}\}.$$

On the other hand $\frac{1}{2}(f_1 + f_2)$ is also a maximizer. If there is a point where f_1, f_2 take different values in \mathcal{V} , then $\frac{1}{2}(f_1 + f_2)$ would take an edge value contradicting Lemma 7.6.15. Therefore $f_1 = f_2$. \square

8 Large deviation principles

Here we put together the results of the previous sections to prove the main results of this paper, namely two versions of a *large deviation principle (LDP)* for fine-mesh limits of random dimer tilings of regions $R \subset \mathbb{R}^3$ with some fixed limiting boundary value b , in the topology induced by the Wasserstein metric on flows d_W introduced in Section 5. Section 1.4 also includes a discussion of our results and a brief description of what a large deviation principle is in general. For more background information, see e.g. [DZ09] or [Var16]. Here we give a slightly more detailed informal description of the main theorems and an outline of the section before getting to formal theorem statements in Sections 8.1 and 8.2. We use results here from throughout the paper, but a lot of the notation in this section was originally introduced in Section 5.

For the large deviation principles, we only work with the boundary flows b which are (i) *boundary asymptotic flows* meaning that b has an extension g to R which is an asymptotic flow (Definition 5.4.8) and (ii) *extendable outside* meaning there exists $\epsilon > 0$ such that b extends to a divergence-free measurable vector field valued in \mathcal{O} on an ϵ neighborhood of R (Definition 8.1.1). Analogous extendability conditions are also required in the large deviation principle for dimer tilings in 2D [CKP01]; see Remark 8.1.2.

In both versions of the LDP we prove, we look at measures supported on dimer tilings of finite regions in $\frac{1}{n}\mathbb{Z}^3$ that cover R (we call these *free-boundary tilings of R at scale n* , see Definition 5.0.1). We can require that the boundary values of these flows converge as $n \rightarrow \infty$ to the fixed boundary value b with either a *soft constraint* or a *hard constraint* on the tilings.

The large deviation principle for dimer tilings in two dimensions [CKP01] uses a hard constraint. In three dimensions, new subtleties arise from the fact that ent can be nonzero on $\partial\mathcal{O}$, and the analogous hard boundary large deviation principle is not true in full generality (see the discussion in Section 1.4 or Example 8.2.6). Instead we prove two versions of an LDP, one with soft boundary constraint and one with hard boundary constraint that holds under an additional condition.

A *soft constraint* means that we choose a sequence of good “thresholds” $(\theta_n)_{n \geq 1}$ with $\theta_n \rightarrow 0$ as $n \rightarrow \infty$, and look at uniform measures ρ_n on free-boundary tilings of R at scale n with boundary values within θ_n of b in the Wasserstein metric $\mathbb{W}_1^{1,1}$ that we use to compare boundary values. The *soft boundary large deviation principle (SB LDP)* says that ρ_n satisfy an LDP, as long as θ_n goes to 0 slowly enough. This is stated precisely in Theorem 8.1.6.

A *hard constraint* means that we choose a sequence of fixed boundary value b_n in the discrete such that b_n converges to b in $\mathbb{W}_1^{1,1}$. We say a boundary value is a *scale n tileable* if there exists a free-boundary tiling τ of R at scale n with that boundary value. If two tilings τ_1, τ_2 have the same boundary values on ∂R , then they are tilings of the same fixed region, so fixing a sequence of scale n tileable boundary values b_n is equivalent to fixing a sequence of *regions* R_n with boundary value b_n . We define $\bar{\rho}_n$ to be uniform measure on free boundary tilings of R at scale n with boundary value b_n , or equivalently as uniform measure on tilings of the fixed region R_n . We show that the measures $(\bar{\rho}_n)_{n \geq 1}$ satisfy an LDP under two conditions: (i) the region R_n is tileable for all n (equivalently, b_n is scale n tileable boundary value) and (ii) the region and boundary value pair (R, b) is *flexible* meaning that for every $x \in \text{Int}(R)$, there exists g extending b and an open set $U \ni x$ such that $\overline{g(U)} \subset \text{Int}(\mathcal{O})$ or equivalently, there exists $f_0 \in AF(R, b)$ such that for every compact set $D \subset \text{Int}(R)$, $\overline{f_0(D)} \subset \text{Int}(\mathcal{O})$ (see

Definition 8.2.1 and Lemma 8.2.2). We call this the *hard boundary large deviation principle* (HB LDP), and it is stated precisely in Theorem 8.2.4.

The condition (R, b) *flexible* is strictly stronger than (R, b) *semi-flexible*, which says that for every point $x \in \text{Int}(R)$, there is an extension g and an open set $U \ni x$ such that $g(U) \subset \mathcal{O} \setminus \mathcal{E}$, or equivalently that b has an extension f_0 which is valued in $\mathcal{O} \setminus \mathcal{E}$ (see Definition 7.6.8 and Lemma 7.6.14). Recall that if (R, b) is not semi-flexible we call it *rigid*.

If (R, b) is *semi-flexible*, then $\text{Ent}(\cdot)$ has a unique maximizer in $AF(R, b)$ (Theorem 7.6.10). Whenever this holds, as a corollary of either LDP we show that “random dimer tilings” of R with boundary values converging to b *concentrate* in the fine-mesh limit on the unique deterministic limiting flow which maximizes $\text{Ent}(\cdot)$ in $AF(R, b)$. This result holds for “random dimer tiling” defined by sampling from any sequence of measures (i.e. ρ_n or $\bar{\rho}_n$) for which an LDP holds, see Corollaries 8.1.9 and 8.2.7.

We summarize the conditions needed for each of the theorems in the following table. Note that in all cases we have the basic assumptions that $R \subset \mathbb{R}^3$ is a compact region which is the closure of a connected domain, ∂R is piecewise smooth, and b is a boundary asymptotic flow which is extendable outside.

(R, b)	SB LDP	Unique Ent maximizer in $AF(R, b)$	HB LDP
rigid	yes	not known in general	no
semi-flexible	yes	yes	no
flexible	yes	yes	yes

We remark that the “no” entries in this table are statements that are provably not true. In particular, there exists (R, b) semi-flexible for which the hard boundary LDP is false; see Example 8.2.6 or the discussion in Section 1.4. See Problem 9.0.7 for discussion of the “not known” entry.

In Section 8.1, we give the precise definitions, conditions, and statement for the soft boundary LDP (Theorem 8.1.6), and in Section 8.2, we do the same for the hard boundary LDP (Theorem 8.2.4), and explain why the hard boundary LDP can be false for (R, b) just semi-flexible (Example 8.2.6). In both cases, we prove *concentration* when (R, b) is semi-flexible and the LDP holds (so in hard boundary case, (R, b) must be flexible) as a corollary (Corollaries 8.1.9 and 8.2.7) and show that proving the LDP is equivalent to proving corresponding upper and lower bounds statements (Theorems 8.1.10 and 8.1.11 for the soft boundary LDP and Theorems 8.2.8 and 8.2.9 for the hard boundary LDP). The rest of the section is dedicated to proving the upper and lower bounds.

The proofs of the lower bounds are somewhat involved. In Section 8.3 we show that if b is extendable outside, then any $g \in AF(R, b)$ can be approximated by a piecewise-constant asymptotic flow on a region slightly larger than R (Proposition 8.3.1). This is where we use the extendable outside condition. Building on this, in Section 8.4 we show that any asymptotic flow can be approximated by the tiling flow of a free-boundary tiling (Theorem 8.4.1). Combined with the patching theorem (Theorem 6.3.5), this is all we need to prove the *soft boundary* lower bound (Theorem 8.1.10), so we prove this in Section 8.5. In Section 8.6 we state and prove a more powerful *generalized patching* theorem (Theorem 8.6.2) and use this to prove the hard boundary lower bound (Theorem 8.2.8).

In Section 8.7 we prove both upper bounds (Theorem 8.1.11 and 8.2.9). To do this, we

prove the soft boundary upper bound (Theorem 8.1.11) and note that this implies the hard boundary upper bound (Theorem 8.2.9).

8.1 Statement and set up: soft boundary LDP

Let $R \subset \mathbb{R}^3$ be a compact region which is the closure of a connected domain, with ∂R piecewise smooth. Recall from Section 5 that for each n , $TF_n(R)$ is the set of all scale n free boundary tiling flows on R . The fine-mesh limits of these with respect to the Wasserstein metric on flows (Theorem 5.3.1) are the asymptotic flows $AF(R)$. The space of asymptotic flows with fixed boundary value b is denoted by $AF(R, b)$. For any compact, piecewise smooth surface $S \subset R$, $T(\cdot, S)$ denotes the *trace operator* which takes an asymptotic or tiling flow to its boundary value on S (see Sections 5.4, 5.5). Recall (Definition 5.4.8) that b is a *boundary asymptotic flow on R* if there exists $g \in AF(R)$ such that $T(g, \partial R) = b$. We restrict our attention to boundary asymptotic flows b which are also *extendable outside*.

Definition 8.1.1. A boundary asymptotic flow b on R is *extendable outside* if there exists $\epsilon > 0$ such that b extends to a divergence-free measurable vector field on an ϵ neighborhood of R .

Remark 8.1.2. The assumption that the boundary asymptotic flow is extendable outside is inherent in [CKP01]. The Lipschitz condition in [CKP01, Theorem 1.1] implies that there is extension of the flow in \mathbb{R}^2 . Such a strong hypothesis is not necessary. However it is easy to build boundary asymptotic flows which are not extendable outside, and some of our current techniques do not work in such cases. Let $R = [-1, 1]^2 \setminus [0, 1]^2$ and consider the flow $f \in AF(R)$ given by

$$f(x) = \begin{cases} (3/4, 0) & \text{if } x \in [-1, 0] \times [0, 1] \\ (0, 3/4) & \text{if } x \in [0, 1] \times [-1, 0] \\ (0, 0) & \text{if } x \in [-1, 0] \times [-1, 0]. \end{cases}$$

Any extension of such a flow close to the origin will have to be valued outside \mathcal{O}_2 by the divergence-free condition. We need b to be extendable outside in our arguments to construct a piecewise-constant approximation \tilde{g} of any flow $g \in AF(R, b)$, where \tilde{g} is supported on a set $\tilde{R} \supset R$ (Proposition 8.3.1). This is an intermediate step in showing that any $g \in AF(R, b)$ can be approximated by a free-boundary tiling $\tau \in T_n(R)$ for n large enough (Theorem 8.4.1). If R is convex, then b is automatically extendable and thus we don't need to add a condition.

The version of the LDP we present in this section has *soft* boundary conditions in the discrete. The sequence of probability measures $(\rho_n)_{n \geq 1}$ which we show satisfy an LDP are uniform probability measures on tiling flows at scale n with boundary values conditioned to be in a sequence of neighborhoods around b which shrink as $n \rightarrow \infty$.

Recall that the metric on boundary values of flows is $\mathbb{W}_1^{1,1}$. To define ρ_n , we first define the following sets.

Definition 8.1.3. Let b be a boundary asymptotic flow and fix a threshold $\theta > 0$. We denote the set of scale n tiling flows on R with boundary values within θ of b by

$$TF_n(R, b, \theta) := \{f_\tau \in TF_n(R) : \mathbb{W}_1^{1,1}(T(f_\tau, \partial R), b) < \theta\}.$$

Note that if θ is too small, $TF_n(R, b, \theta)$ might be empty. However, it will follow from Theorem 8.4.1 that given a fixed θ , if n is large enough then $TF_n(R, b, \theta)$ is nonempty (Corollary 8.4.2).

We say a sequence of thresholds $(\theta_n)_{n \geq 1}$ is **admissible** if $\theta_n \rightarrow 0$ as $n \rightarrow \infty$, but sufficiently slowly so that $TF_n(R, b, \theta_n)$ is nonempty for all n . When the threshold sequence θ_n is understood, we define

$$TF(R, b) := \cup_{n \geq 1} TF_n(R, b, \theta_n).$$

We define a sequence of probability measures ρ_n using an admissible sequence of thresholds.

Definition 8.1.4. For all $n \geq 1$, ρ_n is the uniform probability measure on $TF_n(R, b, \theta_n)$. Further, we define μ_n to be the counting measure on $TF_n(R, b, \theta_n)$ and Z_n to be its partition function, so that $\rho_n = \frac{1}{Z_n} \mu_n$.

Remark 8.1.5. If Unif_n denotes the uniform probability measure on $TF_n(R)$, then ρ_n is the conditional distribution

$$\rho_n(\cdot) = \text{Unif}_n(\cdot \mid D_{b, \theta_n})$$

where D_{b, θ_n} is the event that the boundary value of a flow is within θ_n of b .

Theorem 8.1.6 (Soft boundary large deviation principle). *Let $R \subset \mathbb{R}^3$ be a compact region which the closure of a connected domain, with piecewise smooth boundary ∂R . Let b be a boundary asymptotic flow which is extendable outside.*

There exists a sequence of admissible thresholds $(\theta_n)_{n \geq 1}$ such that the uniform probability measures $(\rho_n)_{n \geq 1}$ on $TF_n(R, b, \theta_n)$ satisfy a large deviation principle in the topology induced by d_W with good rate function $I_b(\cdot)$ and speed $v_n = n^3 \text{Vol}(R)$. Namely for any d_W -Borel measurable set A ,

$$-\inf_{g \in A^\circ} I_b(g) \leq \liminf_{n \rightarrow \infty} v_n^{-1} \log \rho_n(A) \leq \limsup_{n \rightarrow \infty} v_n^{-1} \log \rho_n(A) \leq -\inf_{g \in A} I_b(g) \quad (58)$$

Further, the rate function $I_b(g) = C_b - \text{Ent}(g)$ if g is an asymptotic flow, where $C_b = \max_{f \in AF(R, b)} \text{Ent}(f)$. If g is not an asymptotic flow then $I_b(g) = \infty$.

Remark 8.1.7. The existence of a sequence of thresholds for which the theorem holds follows from Theorem 8.1.10. The only requirement is that $(\theta_n)_{n \geq 0}$ goes to 0 sufficiently slowly.

Remark 8.1.8. The weaker, analogous theorem with free boundary values in the limit would also hold, i.e. there is a large deviation principle for the sequence of uniform measures $(\text{Unif}_n)_{n \geq 1}$ on $TF_n(R)$ from Remark 8.1.5. The rate function in this case is also of the form $C - \text{Ent}(\cdot)$, with $C = \max_{f \in AF(R)} \text{Ent}(f)$.

Under the additional condition that the pair (R, b) is semi-flexible (see Definition 7.6.8), Ent has a unique maximizer in $AF(R, b)$ (Theorem 7.6.10). In this case, Theorem 8.1.6 implies a concentration or weak law of large numbers result for fine-mesh limits of ρ_n -random tiling flows.

Corollary 8.1.9. *Fix $\epsilon > 0$. Assume that (R, b) is semi-flexible so that Ent has a unique maximizer in $AF(R, b)$ which we denote by f_{\max} . Define the event*

$$A_\epsilon = \{f : d_W(f, f_{\max}) > \epsilon\}.$$

Then

$$\rho_n(A_\epsilon) \leq C^{-n^3}$$

where $C > 1$ is a constant depending only on b and R . In other words, for any $\epsilon > 0$, the probability that a tiling flow at scale n sampled from ρ_n (i.e., with boundary value conditioned to be in a shrinking interval around b) differs from the entropy maximizer by more than ϵ goes to 0 exponentially fast as $n \rightarrow \infty$ with rate n^3 .

Proof. Cover $AF(R, b)$ by open neighborhoods B_g around each $g \in AF(R, b)$ so that if $g \neq f_{\max}$ then $\text{Ent}(h) < \text{Ent}(f_{\max})$ for all $h \in \overline{B}_g$, and $B_{f_{\max}}$ is the ϵ -neighborhood of f_{\max} . Since $AF(R, b)$ is compact, this has a finite subcover B_1, \dots, B_k , where B_i is a neighborhood of g_i . Without loss of generality, $B_1 = B_{f_{\max}}$. By Theorem 8.1.6, for n large enough,

$$\rho_n(A_\epsilon) \leq \sum_{i=2}^k \rho_n(B_i) \leq \sum_{i=2}^k \exp(v_n(\text{Ent}(f_i) - \text{Ent}(f_{\max}))),$$

where f_i is the entropy-maximizer in \overline{B}_i . Since $\text{Ent}(f_i) - \text{Ent}(f_{\max}) < 0$ for all $i \neq 1$, this completes the proof. \square

Recall that $\mu_n = Z_n \rho_n$ is counting measure on $TF_n(R, b, \theta_n)$. We define notation for Wasserstein open balls, namely

$$A_\delta(g) = \{h : d_W(h, g) < \delta\}.$$

By [Var16, Lemma 2.3], the large deviation principle for $(\rho_n)_{n \geq 1}$ (Theorem 8.1.6) is implied by local upper and lower bound statements (Theorem 8.1.10 and 8.1.11), plus a property called *exponential tightness*, namely that for any $\alpha < \infty$, there exists a compact set K_α such that, for any closed set C disjoint from K_α ,

$$\limsup_{n \rightarrow \infty} v_n^{-1} \log \rho_n(C) \leq -\alpha. \quad (59)$$

By Corollary 5.4.9, $(AF(R, b), d_W)$ is compact. The space $TF(R, b)$ is countable, and by Theorems 5.3.1 and Theorem 5.5.7, the limit points of $TF(R, b)$ are contained in $AF(R, b)$. Therefore $(AF(R, b) \cup TF(R, b), d_W)$ is compact, from which exponential tightness follows. To prove the soft boundary large deviation principle (Theorem 8.1.6), it remains to prove the following upper and lower bound theorems.

Theorem 8.1.10 (Soft boundary lower bound). *For any $g \in AF(R, b)$,*

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \geq \text{Ent}(g).$$

Theorem 8.1.11 (Soft boundary upper bound). *For any $g \in AF(R, b)$,*

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \leq \text{Ent}(g).$$

8.2 Statement and set up: hard boundary LDP

This section parallels Section 8.1, but the LDP we prove is for measures $(\bar{\rho}_n)_{n \geq 1}$ defined with a *hard boundary constraint* in the discrete, instead of the soft constraint used to define the measures $(\rho_n)_{n \geq 1}$ in Section 8.1.

Again let $R \subset \mathbb{R}^3$ be a compact region which is the closure of a connected domain with ∂R piecewise smooth, and assume that b is a boundary asymptotic flow which is extendable outside. Unlike the soft boundary LDP, we add the condition that the pair (R, b) is *flexible*.

Definition 8.2.1. A pair (R, b) is *flexible* if for all $x \in \text{Int}(R)$, there exists $g \in AF(R, b)$ and an open set $U \ni x$ such that $\overline{g(U)} \subset \text{Int}(\mathcal{O})$.

By completely analogous arguments to the proof of Lemma 7.6.14, we have the following equivalent definition of (R, b) flexible.

Lemma 8.2.2. A pair (R, b) is flexible if and only if there exists $f_0 \in AF(R, b)$ such that for every compact set $D \subset \text{Int}(R)$, $\overline{f_0(D)} \subset \text{Int}(\mathcal{O})$.

Remark 8.2.3. It is not hard to see directly that the flexible definition given in Definition 8.2.1 is satisfied for the 3D regions in the introduction built out of aztec diamonds. On each 2D aztec diamond $R_a = R \cap \{z = a\}$ and each point $x \in R_a$, consider a rectangle inscribed in R_a containing x and with edges parallel to the coordinate axes. Then the flow which is 0 inside the rectangle and linear parallel to the adjacent edge of the rectangle in the four triangles is a 2D asymptotic flow with the right boundary conditions. Averaging these flows for different rectangles gives a flow valued in $\text{Int}(\mathcal{O})$ on R_a , and combining them gives a flow f valued in $\text{Int}(\mathcal{O})$ everywhere in $\text{Int}(R)$ (in fact, f will be valued in the middle slice of $\text{Int}(\mathcal{O})$ where the third coordinate is zero).

A boundary value b_n on ∂R is a *scale n tileable* if there exist a scale n free boundary tiling τ of R such that $T(f_\tau, \partial R) = b_n$. A region $R_n \subset \frac{1}{n}\mathbb{Z}^3$ is a *scale n region with boundary value b_n* if all tilings of R_n have boundary value b_n on ∂R . When b_n is a tileable boundary value, R_n is tileable. Note that implicit in this definition is that all tilings of R_n are scale n free-boundary tilings of R .

Given a sequence of regions R_n with scale n tileable boundary values b_n on ∂R , we define $\bar{\rho}_n$ to be uniform measure on tilings of R_n (equivalently, uniform measure on free-boundary tilings of R with boundary value exactly b_n). If (R, b) is flexible and the boundary values b_n converge to b in $\mathbb{W}_1^{1,1}$ as $n \rightarrow \infty$, we prove that $(\bar{\rho}_n)_{n \geq 1}$ satisfy an LDP.

Theorem 8.2.4 (Hard boundary large deviation principle). *Let $R \subset \mathbb{R}^3$ be a compact region which the closure of a connected domain, with piecewise smooth boundary ∂R . Let b be a boundary asymptotic flow which is extendable outside, and assume that (R, b) is flexible.*

Let $R_n \subset \frac{1}{n}\mathbb{Z}^3$ be a sequence of scale n regions with tileable boundary values b_n converging to b in $\mathbb{W}_1^{1,1}$. Define $\bar{\rho}_n$ to be the uniform probability measure on tilings of R_n .

The measures $(\bar{\rho}_n)_{n \geq 1}$ satisfy a large deviation principle in the topology induced by d_W with good rate function $I_b(\cdot)$ and speed $v_n = n^3 \text{Vol}(R)$. Namely for any d_W -Borel measurable set A ,

$$-\inf_{g \in A^\circ} I_b(g) \leq \liminf_{n \rightarrow \infty} v_n^{-1} \log \bar{\rho}_n(A) \leq \limsup_{n \rightarrow \infty} v_n^{-1} \log \bar{\rho}_n(A) \leq -\inf_{g \in A} I_b(g) \quad (60)$$

Further, the rate function $I_b(g) = C_b - \text{Ent}(g)$ if g is an asymptotic flow, where $C_b = \max_{f \in AF(R, b)} \text{Ent}(f)$. If g is not an asymptotic flow then $I_b(g) = \infty$.

Remark 8.2.5. The large deviation principle in [CKP01] has hard boundary conditions, where the regions R_n approximate R from *within*, i.e. $R_n \subset R$. We instead assume that $R_n \supset R$, and that our regions approximate R from outside.

The *flexible* condition on (R, b) is needed for the *generalized patching theorem* (Theorem 8.6.2). We do not know the exact condition on (R, b) needed for Theorem 8.2.4 to hold, however there do exist regions (R, b) which are just semi-flexible but for which the hard boundary LDP fails.

Example 8.2.6. As discussed in the introduction, there exists semi-flexible region and boundary condition pairs (R, b) with $R \subset \mathbb{R}^3$ for which the hard boundary large deviation principle is false. The region R is a “tilted tube,” and the boundary value b takes values in a face of $\partial\mathcal{O}$. This construction is related to the measures with boundary mean current discussed in Section 4. Recall the definition of the *slabs*

$$L_c = \{(x_1, x_2, x_3) : x_1 + x_2 + x_3 = 2c \text{ or } 2c + 1\}.$$

Each slab is the union of two planes. Any tiling sampled from a measure with mean current $(s_1, s_2, s_3) \in \partial\mathcal{O}$ with $s_1, s_2, s_3 \geq 0$ breaks into a sequence of complete dimer tilings of the slabs (Proposition 4.2.1).

Let $B_n = [-n, n]^3$. Let $A_n(0) = L_0 \cap B_n$, let $A_n(c) = L_c \cap [B_n + (0, 0, 2c)]$, and finally let $A_{n,n} = \cup_{i=-n}^n A_n(i)$. The region R is then defined so that $\frac{1}{n}A_{n,n}$ is a sequence of discrete regions approximating it. We choose the boundary value b to be a constant mean current $s = (s_1, s_2, s_3) \in \partial\mathcal{O}$ with $s_1, s_2, s_3 > 0$. Note that the constant asymptotic flow $g(x) = s \in AF(R, b)$, and by Theorem 4.2.7 has $\text{Ent}(g) = \text{ent}_{\text{loz}}(s) > 0$. Here are two options we could choose for the sequence of discrete regions:

- We define one sequence of regions R_n^1 where for each c such that L_c intersects $A_{n,n}$, $R_n^1 \cap L_c = S_c$ is a region such that a lozenge tiling of S_c has slope s along ∂S_c .
- We define another sequence of regions R_n^2 by alternating between frozen brickwork lozenge tilings. Choose a sequence of ratios (s_1^n, s_2^n, s_3^n) converging to (s_1, s_2, s_3) as $n \rightarrow \infty$. We partition the group of indices c such that L_c intersects $A_{n,n}$, into three groups with sizes proportional to s_1^n, s_2^n, s_3^n . For $i = 1, 2, 3$, for each c in the i^{th} group, we define $R_n^2 \cap L_c$ to be the region tileable by the η_i lozenge brickwork tiling.

By results for 2D lozenge tilings and the relationships established in Section 4, the hard boundary LDP would hold for the sequence R_n^1 . On the other hand, R_n^2 is frozen for all n , so the number of free boundary tilings of R_n^2 is 1 for all n . While this unique tiling does approximate the constant flow $g(x) = s$, the corresponding lower bound for the LDP does not hold.

More generally, one might conjecture that the hard boundary LDP fails when there exist regions in $\text{Int}(R)$ where the Ent-maximizing flow is valued in the faces of $\partial\mathcal{O}$. However we do not know if this is a necessary or sufficient condition for the hard boundary LDP to fail, or if there exists a region $R \subset \mathbb{R}^3$ where the Ent-maximizing flow takes face values in the interior of R but not all of R . See Problem 9.0.8.

The analogous concentration or weak law of large numbers result that held for the soft boundary measures ρ_n (Corollary 8.1.9) also holds for the hard boundary measures $\bar{\rho}_n$. Note that flexible implies semi-flexible, so the maximizer $f_{\text{max}} \in AF(R, b)$ is unique by Theorem 7.6.10.

Corollary 8.2.7. *Assume that (R, b) is flexible and $\bar{\rho}_n$ are as in Theorem 8.2.4. Let f_{\max} denote the unique maximizer of Ent in $AF(R, b)$. Define the event*

$$A_\epsilon = \{f : d_W(f, f_{\max}) > \epsilon\}.$$

Then

$$\bar{\rho}_n(A_\epsilon) \leq C^{-n^3}$$

where $C > 1$ is a constant depending only on R and b . In other words, for any $\epsilon > 0$, the probability that a tiling flow at scale n sampled from $\bar{\rho}_n$ (i.e., a tiling of R_n) differs from the entropy maximizer by more than ϵ goes to 0 exponentially fast as $n \rightarrow \infty$ with rate n^3 .

Proof. Analogous to the proof of Corollary 8.1.9 with ρ_n replaced by $\bar{\rho}_n$ and Theorem 8.1.6 replaced by Theorem 8.2.4. \square

Like Theorem 8.1.6, by [Var16, Lemma 2.3], to prove Theorem 8.2.4 it suffices to show that the measures $(\bar{\rho}_n)_{n \geq 1}$ satisfy local upper and lower bound statements (Theorems 8.2.8 and 8.2.9) plus the *exponential tightness property* stated for ρ_n in Equation (59), which follows by analogous straightforward arguments for $\bar{\rho}_n$.

We let \bar{Z}_n denote the partition function of $\bar{\rho}_n$ and $\bar{\mu}_n = \bar{Z}_n \bar{\rho}_n$ the corresponding counting measures.

Theorem 8.2.8 (Hard boundary lower bound). *For any $g \in AF(R, b)$,*

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} v_n^{-1} \log \bar{\mu}_n(A_\delta(g)) \geq Ent(g).$$

Theorem 8.2.9 (Hard boundary upper bound). *For any $g \in AF(R, b)$,*

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} v_n^{-1} \log \bar{\mu}_n(A_\delta(g)) \leq Ent(g).$$

8.3 Piecewise constant approximation

The goal of this section is to show that if b is extendable outside, then any $g \in AF(R, b)$ is well-approximated in the Wasserstein metric on flows by an asymptotic flow \tilde{g} which is piecewise-constant, taking constant values on a mesh \mathcal{X} of small tetrahedra covering R (see Remark 8.3.3 for why tetrahedra).

Proposition 8.3.1. *Fix $\epsilon > 0$, and suppose that b is a boundary asymptotic flow which is extendable outside. For any $g \in AF(R, b)$, there exists $\delta > 0$ and a δ -mesh of tetrahedra \mathcal{X} covering R with the following properties. Let $\tilde{R} = \cup_{X \in \mathcal{X}} X$. There exists a flow \tilde{g} satisfying:*

- $\tilde{g} \in AF(\tilde{R})$;
- $d_W(g, \tilde{g}) < \epsilon$;
- For each $X \in \mathcal{X}$, $\tilde{g}|_X = \tilde{g}_X$ is constant;
- \tilde{g} is valued strictly in $Int(\mathcal{O})$;
- \tilde{g} takes only rational values.




Figure 34: A cube cut into one regular tetrahedron and four right-angled tetrahedra. The second picture shows the same tetrahedra moved apart.

Remark 8.3.2. We need the condition that b is extendable outside so that we can take \tilde{R} to contain R . If b is not extendable outside, a similar construction works, but the resulting piecewise-constant flow will be an asymptotic flow on a region R' slightly smaller than R instead.

Remark 8.3.3. Tetrahedral mesh. The fact that the mesh in this construction is built out of tetrahedra is necessary to ensure that \tilde{g} is divergence-free (needed for \tilde{g} to be an asymptotic flow). This is because a divergence-free flow on a polyhedron with F faces is determined by its flow through $F - 1$ of them. Since we have 3 free parameters to specify \tilde{g} on one polyhedron, we need $F - 1 \leq 3$. The only polyhedra that satisfy this are tetrahedra.

However, regular tetrahedra alone do not tile 3-space¹, so we cannot take all elements of the mesh to be identical. Instead, 3-space can be tiled by regular tetrahedra and right-angled tetrahedra. To see this, note that cubes tile 3-space, and a cube can be cut into four right-angled tetrahedra and one regular tetrahedron (see Figure 34). The faces of the regular tetrahedron have normal vectors of the form $(\pm 1, \pm 1, \pm 1)$, while the right-angled tetrahedra have four coordinate plane faces and one face with normal vector $(\pm 1, \pm 1, \pm 1)$. For technical reasons (see the proof of Theorem 8.4.1, in particular Lemma 8.4.6), our arguments are simplified by assuming that the faces of the tetrahedra are always contained in one of these two types of planes. In the proof of Proposition 8.3.1 we will also use this to say that the possible normal vectors to the tetrahedra can be assumed to form a finite set. We assume throughout that our tetrahedral mesh is built out of regular and right-angled tetrahedra.

Proof of Proposition 8.3.1. Since b is extendable outside, there exist $\alpha_0 > 0$ such that g can be extended to $g' \in AF(R^{\alpha_0})$, where R^{α_0} is

$$R^{\alpha_0} = \{x \in \mathbb{R}^3 : d(x, R) \leq \alpha_0\}.$$

Given this, for any $0 < \alpha < \alpha_0$, we can approximate g by a *continuous* asymptotic flow $g_\alpha \in AF(R^{\alpha_0 - \alpha})$:

$$g_\alpha(x) := \frac{1}{|B_\alpha(x)|} \int_{B_\alpha(x)} g'(y) \, dy, \quad x \in R^{\alpha_0 - \alpha}.$$

¹Over 2,000 years ago, Aristotle (mistakenly) claimed in *De Caelo*, Book III Part 8 [AriBC] that regular tetrahedra do tile 3-space. It took around 1,000 years for the mistake to be fixed, see [LZ12] for a detailed account of the story.

As $\alpha \rightarrow 0$, $d_W(g, g_\alpha |_R) \rightarrow 0$. We construct a piecewise-constant, divergence-free approximation u of g_α , then modify it to construct \tilde{g} also satisfying the last two conditions.

For any fixed $\delta < \alpha_0 - \alpha$ (to be specified more precisely later), we take a δ -tetrahedral mesh \mathcal{X} built from regular and right-angled tetrahedra (see Remark 8.3.3) such that $X \cap R \neq \emptyset$ for all $X \in \mathcal{X}$. Let $\tilde{R} = \cup_{X \in \mathcal{X}} X$ and note that $R \subset \tilde{R} \subset R^{\alpha_0 - \alpha}$.

Consider one tetrahedron $X \in \mathcal{X}$. Let $\zeta_1, \zeta_2, \zeta_3, \zeta_4$ denote the faces of X and let n_1, n_2, n_3, n_4 denote their outward pointing normal vectors. Define a vector u_X by

$$u_X \cdot n_i = \frac{1}{\text{area}(\zeta_i)} \int_{\zeta_i} \langle g_\alpha, n_i \rangle dA \quad i = 1, 2, 3.$$

Since g_α is divergence-free on X ,

$$u_X \cdot n_4 = \frac{1}{\text{area}(\zeta_4)} \int_{\zeta_4} \langle g_\alpha, n_4 \rangle dA.$$

Define $u(x) := u_X$ for $x \in X$.

It remains to show that (up to multiplying by a constant $\lambda \leq 1$ but very close to 1) $\lambda u \in AF(\tilde{R})$ and bound $d_W(\lambda u, g_\alpha |_{\tilde{R}})$.

Since g_α is continuous and $R^{\alpha_0 - \alpha}$ is compact, g_α is uniformly continuous on $R^{\alpha_0 - \alpha}$. Thus given any $\beta > 0$ there exists $\theta > 0$ such that $|x - y| < \theta$ implies $|g_\alpha(x) - g_\alpha(y)| < \beta$.

Fixing β , we now require that $\delta < \theta$ so that uniform continuity implies that for any $X \in \mathcal{X}$ and point $x \in X$, we have that $|g_\alpha(x) - \text{avg}_X g_\alpha| < \beta$. The normal vectors n_1, n_2, n_3 to three faces of X are linearly independent but not necessarily orthogonal. However since all $X \in \mathcal{X}$ are of one of five forms (see Figure 34), there is a constant $K > 0$ independent of $X \in \mathcal{X}$ so that

$$|u_X - \text{avg}_X g_\alpha| \leq K \sum_{i=1}^3 |\text{avg}_{\zeta_i} (g_\alpha \cdot n_i) - \text{avg}_X (g_\alpha \cdot n_i)| < 3K\beta.$$

Therefore for all $x \in \tilde{R}$,

$$|u(x) - g_\alpha(x)| < (3K + 1)\beta. \quad (61)$$

Replacing u by λu with $\lambda = 1 - (3K + 1)\beta - \beta$, the new flow $\lambda u \in AF(\tilde{R})$ and in fact is valued in $\text{Int}(\mathcal{O})$. Further by Proposition 5.2.11, there is another constant $C > 0$ such that

$$d_W(\lambda u, g_\alpha |_{\tilde{R}}) < C(6K + 3)\beta. \quad (62)$$

By the triangle inequality,

$$d_W(\lambda u, g) \leq d_W(\lambda u, g_\alpha |_{\tilde{R}}) + d_W(g_\alpha |_{\tilde{R} \setminus R}, 0) + d_W(g, g_\alpha |_R).$$

The first term is controlled by Equation (62). The second is bounded by a fixed constant times δ , and the third is bounded by a fixed constant times α . Therefore taking $\alpha, \beta > 0$ small enough and correspondingly taking $\delta < \max\{\alpha_0 - \alpha, \theta\}$, $d_W(\lambda u, g)$ can be made arbitrarily small.

The flow $\lambda u \in AF(\tilde{R})$ and is valued in $\text{Int}(\mathcal{O})$. Finally we modify λu as follows to construct \tilde{g} which also takes rational values. To do this, we solve the linear constraint problem to make the values of the flow rational without breaking the divergence-free condition.

Let M be the number of tetrahedra in the mesh \mathcal{X} . Enumerate the faces of the tetrahedra by a_1, \dots, a_m . Choose a unit normal vector n_i for each face. For any flow f , let $F(f) = (F_1(f), \dots, F_m(f))$, where $F_i(f) = \int_{a_i} \langle f, n_i \rangle dx$. Note that if v is a piecewise-constant flow on the mesh, then $F(v)$ determines v .

If $F(v)$ corresponds to a divergence-free piecewise-constant flow v , then it satisfies a matrix A of M linear constraints of the form

$$\pm F_{k_1}(v) \pm F_{k_2}(v) \pm F_{k_3}(v) \pm F_{k_4}(v) = 0$$

for a_{k_j} , $j \in \{1, 2, 3, 4\}$ the faces of a tetrahedron $X \in \mathcal{X}$ (the signs are determined by the normal vector orientation, the four terms should all be for flow oriented out of X). Thus $F(v)$ solves $AF(v) = 0$. Since A has integer entries, there is a rational basis for the space of solutions Y of $AY = 0$. Any other solution can be written as a linear combination of the rational ones, so rational solutions are dense.

Thus we can find \tilde{g} such that \tilde{g}_X takes all rational values and $|\tilde{g}_X - (1 - \delta_1)u_X|$ is as small as needed. Applying Proposition 5.2.11 again completes the proof. \square

8.4 Existence of tiling approximations

Building on the approximation result in the previous section, we now show that if b is extendable outside then any $g \in AF(R, b)$ can be approximated in Wasserstein distance by a tiling flow. More precisely:

Theorem 8.4.1. *Fix $\delta > 0$ and suppose b is a boundary asymptotic flow which is extendable outside. For any $g \in AF(R, b)$, there exists $n(\delta)$ such that if $n \geq n(\delta)$, then there is a free boundary tiling $\tau \in T_n(R)$ such that $f_\tau \in A_\delta(g)$.*

The two dimensional analog of this theorem (i.e. [CKP01, Prop.3.2]) is the statement that any asymptotic height function can be approximated by the height function of a tiling. In particular, one can choose the maximal height function (analog of f_τ) less than the given asymptotic height function (analog of g). There is no analogous notion of “maximal” tiling flow, so our argument in three dimensions is more complicated, and relies on an explicit construction.

We call the explicit construction in the proof of Theorem 8.4.1 the “shinning light construction.” The first step is to build piecewise-linear “channels.” We give a method for tiling the channels and show that we can glue them together to construct a tiling of the whole region. The channels are tubular neighborhoods of the flow lines of a tiling flow approximating a piecewise-constant flow as constructed in Proposition 8.3.1. We call it the “shining light construction” because we imagine the flow as beams of light bending through the channels.

Before proving Theorem 8.4.1, we note that the existence of an admissible sequences of thresholds $(\theta_n)_{n \geq 1}$ follows as a straightforward corollary.

Corollary 8.4.2. *For any boundary asymptotic flow b which is extendable outside and any threshold $\theta > 0$, $TF_n(R, b, \theta)$ is nonempty for n large enough. In particular, admissible sequences of thresholds $(\theta_n)_{n \geq 1}$ exist for any boundary asymptotic flow b which is extendable outside.*

Proof. Recall that $T(\cdot, \partial R) : AF(R) \rightarrow \mathcal{M}^s(R)$ is the boundary value operator and choose $g \in T^{-1}(b)$. By Theorem 8.4.1, for any $n \geq n(\delta)$ there exists a tiling $\tau \in T_n(R)$ with $d_W(f_\tau, g) < \delta$. Since T is uniformly continuous (Theorem 5.5.7), we can choose $\delta > 0$ so that $d_W(g, f_\tau) < \delta$ implies $d_W(b, T(f_\tau)) < \theta$. \square

We now proceed to the explicit construction. Recall that η_i is the i^{th} positively-oriented unit coordinate vector and that e_i denotes the edge in \mathbb{Z}^3 connecting the origin to η_i . Similarly, $-e_i$ is the edge connecting the origin to $-\eta_i$.

Let τ_1 denote the brickwork tiling where all tiles are $-\eta_1$ bricks. To prove Theorem 8.4.1, we show that we can construct a tiling τ so that the flow corresponding to the double dimer tiling (τ, τ_1) is close to the flow $g + \eta_1$. A double dimer tiling consists of a collection of oriented infinite paths, finite loops, and double edges. See Section 2.1 and Section 7.3.

Since τ_1 consists of only $-\eta_1$ tiles, for any other tiling τ , (τ, τ_1) consists of only infinite paths and double edges (i.e. no finite loops). The double dimer flow $f_{(\tau, \tau_1)} = f_\tau - f_{\tau_1}$ is 0 whenever the tilings agree, and otherwise points in the direction of the oriented infinite path.

For $x \in \mathbb{Z}^3$, let $\tau(x)$ denote the tile at x in τ . We say that a tiling τ of \mathbb{Z}^3 is *periodic* if there exist even integers $r_1, r_2, r_3 > 0$ such that $\tau(x)$ is equal to its translates $\tau(x + r_1\eta_1) = \tau(x + r_2\eta_2) = \tau(x + r_3\eta_3)$ for all $x \in \mathbb{Z}^3$. For periodic tilings, we can define a notion of the *mean current of a tiling*, denoted $s(\tau)$, as the average direction of the tiles in any $r_1 \times r_2 \times r_3$ box.

We give a method for constructing a periodic tiling τ_v of \mathbb{Z}^3 of a fixed, rational mean current $v \in \mathcal{O}$. This construction will serve as a building block in the proof of Theorem 8.4.1.

Construction of tiling τ_v .

First we give a construction for $v = (v_1, v_2, v_3) \in \partial\mathcal{O} \cap \mathbb{Q}^3$ with $v_1 \geq 0$, then we adapt this to the general case.

Here we view dimer tiles a in a tiling as vectors directed from even to odd. When we subtract a tiling, we reverse the direction of its dimers. Since v is rational and has nonzero norm, we can find a sequence of tiles $a_1, \dots, a_r \in \{\eta_1, \text{sign}(v_2)\eta_2, \text{sign}(v_3)\eta_3\}$ such that $a_1 + \dots + a_r + r\eta_1$ is parallel to $v + \eta_1$.

Below by a *plane with normal vector $(1, 1, 1)$* , we mean a collection of cubes in \mathbb{Z}^3 with coordinates $\{(x_1, x_2, x_3) : x_1 + x_2 + x_3 = c\}$ for some constant $c \in \mathbb{Z}$. We analogously define a plane with normal vector $(\pm 1, \pm 1, \pm 1)$ to be the modification of this with appropriate signs.

Choose a plane C_0 with normal vector $\xi = (1, \text{sign}(v_2), \text{sign}(v_3))$. Let C_k denote $C_0 + (0, 0, k)$ for all $k \in \mathbb{Z}$. Further, assume that the cubes on C_0 are even, so that edges parallel to one of $\{\eta_1, \text{sign}(v_2)\eta_2, \text{sign}(v_3)\eta_3\}$ connect cubes on C_0 to cubes on C_1 . Since $v \in \partial\mathcal{O}$, any tiling with mean current v splits into perfect dimer tilings of *slabs* L_k which consist of unions of adjacent planes $L_k := C_{2k} \cup C_{2k+1}$, $k \in \mathbb{Z}$ (see Section 4). By Proposition 4.2.1, each slab is a copy of the 2-dimensional hexagonal lattice. There are three 3D dimer tilings of $C_{2k} \cup C_{2k+1}$ consisting of only one type of dimer, and these correspond to the three brickwork lozenge tilings using one type of lozenge. (See Figure 19 for a review of the correspondence between 3D dimers and lozenges.)

Restricted to each slab L_k , τ_v will be one of the three brickwork lozenge tilings. On $L_0 = C_0 \cup C_1$, τ_v will be the a_1 brickwork lozenge tiling. On $L_1 = C_2 \cup C_3$, τ_v will be the a_2 brickwork lozenge tiling. We continue this by repeating the periodic sequence a_1, \dots, a_r forwards and




Figure 35: Above is an example of $v = (v_1, v_2, v_3)$ and its relationship to $w(v) = w$.

backwards in k to choose the tile type for τ_v on all other slabs $L_k = C_{2k} \cup C_{2k+1}$.

The reference tiling τ_1 consists of all $-\eta_1$ tiles, which connect C_{2k} to C_{2k-1} . Subtracting τ_1 , the tiles in $-\tau_1$ connect C_{2k-1} to C_{2k} , meaning that they connected the “odd” half of L_{k-1} to the “even” half of L_k . Hence in the double dimer tiling (τ_v, τ_1) , every tile is on an infinite path. Along each infinite path, (τ_v, τ_1) consists of the periodic sequence of tiles parallel to $\dots a_1, \eta_1, a_2, \eta_1, \dots, \eta_1, a_r, \dots$. In particular, all infinite paths are parallel to $v + \eta_1$. This completes the construction for $v \in \partial\mathcal{O} \cap \mathbb{Q}^3$, $v_1 \geq 0$.

Now we extend the construction to any $v \in \mathcal{O} \cap \mathbb{Q}^3$, $v \neq -\eta_1$. Let p_v be the line through $-\eta_1 = (-1, 0, 0)$ and $v = (v_1, v_2, v_3)$. Let $w = w(v)$ be the intersection of p_v with $\{u = (u_1, u_2, u_3) \in \partial\mathcal{O} : u_1 \geq 0\}$. See Figure 35. The relationship between v, w will be sufficiently important that we record it as a definition.

Definition 8.4.3. Fix $v \in \mathcal{O}$, $v \neq -\eta_1$, and let p_v be the line through $-\eta_1$ and v . We define $w(v)$ to be the intersection of p_v with the part of $\partial\mathcal{O}$ with non-negative first coordinate.

Note that if v is rational, $w = w(v) \in \partial\mathcal{O}$ is rational. Since the first coordinate of w is non-negative, we can construct τ_w as described above. In (τ_w, τ_1) , every tile is along an infinite path. On the other hand, in (τ_1, τ_1) none of the tiles are along infinite paths. To construct (τ_v, τ_1) , we interpolate between these two options by choosing an intermediate density of infinite paths.

If the line $p_v(t)$ is parameterized so that $p_v(1) = (-1, 0, 0)$ and $p_v(0) = w$, let $a \in [0, 1]$ be such that $p_v(a) = v$. Since v rational, w and a are also rational. Thus there exist periodic patterns of cubes in C_0 with density a . To construct (τ_v, τ_1) , we fix a choice of periodic pattern of cubes on C_0 with density a . We delete all the infinite paths in (τ_w, τ_1) which do not go through one of the chosen cubes on C_0 and replace them with tiles parallel to $-\eta_1$. The resulting tiling is τ_v .

The tilings τ_v have a few important properties which we highlight.

Lemma 8.4.4. *Let $v \in \mathcal{O}; v \neq -\eta_1$.*

1. τ_v has mean current v ;
2. Let $w = w(v)$ be as in Definition 8.4.3. The infinite paths in τ_v are parallel to $w + \eta_1$.
3. For any rational plane P , the restriction of τ_v to P is doubly periodic, with period depending on r (the number of tiles a_1, \dots, a_r used to approximate $w(v)$), the choice of periodic pattern of cubes in C_0 and P .

Remark 8.4.5. Note that τ_v is not uniquely determined by v . It depends on the sequence of tiles a_1, \dots, a_r used to approximate v , and on and periodic pattern of initial sites on C_0 .

We now show that pieces of τ_v, τ_u can “glued” along a plane P , as long as v, u have the same flow through P . The amount of space k we need to glue is a constant depending only on the period of the tilings τ_u and τ_v .

Lemma 8.4.6. *Suppose that $u, v \in \mathcal{O} \cap \mathbb{Q}^3; u, v \neq \eta_1$ and τ_u, τ_v are tilings as in Lemma 8.4.4. Suppose that P is a coordinate plane or plane with normal vector of the form $(\pm 1, \pm 1, \pm 1)$. In both cases we denote the normal vector by n_P . Let r be such that τ_u and τ_v are periodic in P with fundamental domain an $r \times r$ parallelogram. If $v \cdot n_P = u \cdot n_P$, then there is an even integer $k > 0$ (depending on r and P) such that τ_v restricted to the left half-space of P can be connected to τ_u restricted to the right half-space of $P + k$ for some $k = O(r)$. Further, the connecting tiling τ is also periodic in P with fundamental domain an $r \times r$ parallelogram.*

Remark 8.4.7. We restrict to these two types of planes P since the the faces of tetrahedra in the mesh used to define the piecewise constant approximation (Proposition 8.3.1) are contained in one of these two types of planes; see Remark 8.3.3. The analogous result for other planes also holds, but we do not need it.

Proof. By Lemma 8.4.4, τ_u, τ_v are periodic, so there exists $r > 0$ finite and determined by τ_u, τ_v, P such that τ_u is periodic on P with fundamental domain $R_0 \subset P$, where R_0 an $r \times r$ parallelogram contained in P , and similarly τ_v is periodic on $P + k$ with fundamental domain also an $r \times r$ parallelogram in $P + k$. Let R be the parallelopiped region parallel to n_P between one fundamental domain $R_0 \subset P$ and another $R_k \subset P + k$.

Let R/\sim be R with opposite faces other than R_0 and R_k paired (i.e., R/\sim is a 2-dimensional torus crossed with an interval). Given the periodicity of τ_u, τ_v , to that show the region between P and $P + k$ is tileable with $\tau_u|_P$ and $\tau_v|_{P+k}$, it suffices to show that R/\sim is tileable with $\tau_u|_{R_0}$ and $\tau_v|_{R_k}$.

To show that this region is tileable we use the same techniques as in Section 6. In other words, first we show that R/\sim with $\tau_u|_{R_0}$ and $\tau_v|_{R_k}$ is balanced, and then use Hall’s matching theorem (6.1.3). The setting here is more elementary than what we consider in Section 6, since here the tilings defining the boundary conditions are completely periodic.

Since k is even, any perpendicular slice of R/\sim is a fundamental domain for τ_u or τ_v , the condition $v \cdot n_P = u \cdot n_P$ is equivalent to

$$\sum_{e \text{ intersecting } R_0} v_{\tau_u}(e) \cdot n_P = \sum_{e \text{ intersecting } R_k} v_{\tau_v}(e) \cdot n_P. \quad (63)$$

Equation (63) is in turn equivalent to R/\sim with boundary conditions $\tau_u|_{R_0}$ and $\tau_v|_{R_k}$ being balanced.

Since the region is balanced, by Hall's matching theorem (Theorem 6.1.3) it is not tileable if and only if there exists a counterexample region U which is a strict subset. Since U is a strict subset, U has a nonempty interior boundary $S \subset \partial U$. Let $T = \partial U \setminus S$. By Proposition 6.2.3,

$$\text{imbalance}(U) = \frac{1}{6} \left(\text{white}(T) - \text{black}(T) - \text{area}(S) \right).$$

Since $\text{area}(T) \leq 2r^2$, $\text{white}(T) \leq 2r^2$. Since the region is tileable with boundary condition from just one of the tilings, S must connect R_0 and R_k , if $k > r$, by Proposition 6.5.1, there is a universal constant κ such that $\text{area}(S) \geq \kappa kr$. Therefore by Proposition 6.2.3,

$$\text{imbalance}(U) \leq \frac{2r^2 - \kappa kr}{6}.$$

Choosing $k = cr$ for some constant $c > 2/\kappa$, U is not a counterexample which contradicts the assumption that the region is not tileable. This completes the proof. \square

Using the tilings τ_v as our building blocks and their gluing properties to put them together, we now proceed to prove Theorem 8.4.1.

Proof of Theorem 8.4.1. Choose a scale $\epsilon > 0$ so that the piecewise-constant approximation \tilde{g} from Proposition 8.3.1 on an ϵ -scale tetrahedral mesh $\mathcal{X} = \{X_1, \dots, X_M\}$ satisfies $d_W(g, \tilde{g}) < \delta/2$ and hence $A_{\delta/2}(\tilde{g}) \subset A_\delta(g)$. We assume that all $X \in \mathcal{X}$ are regular or right-angled tetrahedra so that all their faces are contained in coordinate planes or planes with normal vector $(\pm 1, \pm 1, \pm 1)$. Recall that $\tilde{g} \in AF(\tilde{R})$ and that $R \subset \tilde{R}$, so any free boundary tiling of \tilde{R} can be restricted to a free boundary tiling of R . To prove the theorem, it suffices to construct $\tau \in T_n(\tilde{R})$ with $d_W(f_\tau, \tilde{g}) < \delta/2$.

Constructing channels. We construct *channels* C_1, \dots, C_K which are disjoint, partition \tilde{R} and will be nicely chosen tubular neighborhoods of a modification of the flow lines of $\tilde{g} + \eta_1$. For any $X_j \in \mathcal{X}$, let $v_j := \tilde{g}|_{X_j}$. Recall Definition 8.4.3, which relates a vector v with $w(v)$, which is the direction of the infinite paths in a periodic tiling τ_v . For each channel C_i , the intersection $C_i \cap X_j$ will be a tube parallel to

$$w(v_j) + \eta_1.$$

Since \tilde{g} is valued strictly in $\text{Int}(\mathcal{O})$, $v_j \neq -\eta_1$ for all $X_j \in \mathcal{X}$, and hence $w(v_j)$ is well-defined everywhere. As shorthand, we let $w(\tilde{g})$ be the piecewise-constant flow equal to $w(v_j)$ on X_j . The definitions are made so that if τ_{v_j} is a periodic tiling built by the construction earlier in this section, the infinite paths in (τ_{v_j}, τ_1) move parallel to the direction of the channel on X_j . The values of \tilde{g} change on the boundaries ∂X_j of tetrahedra in the mesh. We choose the channels C_i to be thin enough as follows so that, viewing C_i as a sequence of open tubes, each end of the tube $C_i \cap X_j$ is contained in a single plane (i.e., each end is contained in a single face of ∂X_j).

Project the corners and edges of X_j onto ∂X_j along $w(v_j) + \eta_1$. Call these projections γ_j . The points $\gamma_j \subset \partial X_j$ are the ones along a flow line of $w(\tilde{g}) + \eta_1$ that goes through an edge of X_j . The lines γ_j divide the faces of X_j into between 1 and 3 sections.

We further divide X_j by taking into account the iterated projections of the corners and edges of all the other tetrahedra in the mesh. In other words for all j , if X_k is a neighbor of X_j , then we project γ_j onto ∂X_k by orthogonal projection along $w(v_k) + \eta_1$. We iterate this for




Figure 36: On the left is a face of one tetrahedron. The segments are the ends of channels, the smaller blue regions are places where we do not delete infinite paths. The width of the white area is $O(r/n)$. The figure on the right is a 2D schematic showing two channels C, C' , with the tubes T_j and connector regions B_j labeled along C . The width of the white area between C and C' is $O(r/n)$.

all tetrahedra until there is a projection of γ_j on ∂X_k for all $\{k, j\}$ pairs. See the left figure in Figure 36.

The result is that for each $j \in \{1, \dots, M\}$, each triangular face of ∂X_j is partitioned into between 1 and 3^M pieces, and X_j is partitioned into tubes parallel to $w(v_j) + \eta_1$ with these pieces as their ends. Each sequence of successive tubes glued on their intersections with ∂X_j is a *channel*. The collection of channels C_1, \dots, C_K is pairwise disjoint and covers \tilde{R} . Since $w(\tilde{g}) + \eta_1$ has positive first coordinate everywhere, each channel connects a patch on $\partial \tilde{R}$ to another.

Tiling a channel. Fix n large and a choice of channel C . We construct a scale n tiling of C which has only $-\eta_1$ tiles in a neighborhood of ∂C of constant-order width in n , and use this to say that we can put together the tilings of the channels together to construct one tiling of the whole region.

Let T_1, \dots, T_m be the sequence of tubes of the form $X_j \cap C$ in order from one intersection of C with $\partial \tilde{R}$ to the other. Let v_1, \dots, v_m be the corresponding values of \tilde{g} on the tubes. Consider the \mathbb{Z}^3 tilings $\tau_{v_1}, \dots, \tau_{v_m}$ constructed earlier in this section using the reference tiling τ_1 . Recall that for each v_i , all tiles in (τ_{v_i}, τ_1) are either double tiles (which must be $-\eta_1$ tiles) or are on an infinite path, and that all infinite paths follow the same periodic sequence. Let r_j be the period of τ_{v_j} , i.e. if $\tau(x)$ denotes the tile at x in τ , then for $j = 1, \dots, m$, r_j is such that translates $\tau_{v_j}(x + r_j \eta) = \tau_{v_j}(x)$ for all unit vectors η .

The main operation we will use is that for any infinite path $\ell \subset (\tau_{v_j}, \tau_1)$, we can modify τ_{v_j} by “shifting” all the tiles along ℓ , i.e. by removing all the tiles on τ_{v_j} along ℓ and replacing them with $-\eta_1$ tiles. We refer to this as *deleting* the path ℓ . The idea is to delete paths that would exit the channel before hitting $\partial \tilde{R}$, and then to bound the number of paths that we delete.

Let $\pi_j \subset \partial X_j$ be the starting end of T_j , so that T_j is a tube connecting π_j to π_{j+1} . For each j , we start by restricting τ_{v_j} to T_j . Any infinite path $\ell \subset (\tau_{v_j}, \tau_1)$ which enters T_j in π_j must exit through $\partial T_j \setminus \pi_j$, since paths in (τ_{v_j}, τ_1) always have direction with positive η_1 component.

First, we delete all infinite paths $\ell \subset (\tau_{v_j}, \tau_1)$ which do not enter T_j through π_j and exit for

the first time through π_{j+1} , replacing the tiles of τ_{v_j} along these paths with $-\eta_1$ tiles. Note that this includes deleting all infinite paths which do not intersect T_j .

By Lemma 8.4.4, the infinite paths in (τ_{v_j}, τ_1) have asymptotic direction $w(v_j)$, with oscillation bounded by the length of the periodic sequence used to construct τ_{v_j} , which is $O(r_j)$. Since the direction of the tube T_j is also $w(v_j)$, any infinite path in (τ_{v_j}, τ_1) that enters T_j through π_j and exits through $T_j \setminus \pi_{j+1}$ is within $O(r_j)$ distance counted as number of edges in $\frac{1}{n}\mathbb{Z}^3$ of ∂C along T_j . Similarly, any path which enters T_j through $\partial T_j \setminus \pi_j$ would also remain within $O(r_j)$ distance in number of $\frac{1}{n}\mathbb{Z}^3$ edges of ∂C along T_j . In summary, the paths that we delete which intersect T_j are all contained in an neighborhood of $\partial C \cap T_j$ of width $O(r)$ in edge distance in $\frac{1}{n}\mathbb{Z}^3$, corresponding to a neighborhood of Euclidean width $O(r/n)$ (recall that r is a constant independent of n).

Second (to avoid issues with corners and edges of tetrahedra, and to isolate channels from each other), we delete all infinite paths which are within a Euclidean neighborhood of width $1000/n$ of $\partial T_j \setminus \{\pi_j \cup \pi_{j+1}\}$ (i.e., 1000 lattice cubes in $\frac{1}{n}\mathbb{Z}^3$). By the same logic as above, these are still contained in an $O(r/n)$ -width neighborhood of ∂C . We call this tiling τ'_{v_j} .

Third, we want to glue the tiling on T_j to the tiling on T_{j+1} . To do this, we cut out a neighborhood of width $O(r) = O(r_1, \dots, r_m)$ in $\frac{1}{n}\mathbb{Z}^3$ lattice cubes around π_{j+1} (the face shared by T_j and T_{j+1}) which we call the *connector region* B_j . Let α_j, α_{j+1} be the ends of the connector region (i.e. translates of π_{j+1}), see the purple region in right side figure in Figure 36. Let P_j be the plane containing α_j and P_{j+1} be the plane containing α_{j+1} .

Since τ_{v_j} and $\tau_{v_{j+1}}$ are periodic tilings of periods r_j, r_{j+1} , and since B_j has width $O(r)$, by Lemma 8.4.6 we can construct a tiling σ_j of \mathbb{Z}^3 such that it agrees with τ_{v_j} in the left half-space of P_j and $\tau_{v_{j+1}}$ on the right half-space of P_{j+1} , and fills in the region in-between in a periodic way with period $O(r)$. We can do this for all $j = 1, \dots, m$.

Overlaying (σ_j, τ_1) , we again get a collection of infinite paths and double tiles. First, we delete all infinite paths in σ_j which were deleted to construct τ'_{v_j} and $\tau'_{v_{j+1}}$ from τ_{v_j} and $\tau_{v_{j+1}}$. Second, we delete any infinite paths which exit C between α_j and α_{j+1} (i.e., any paths which exits C along the connector B_j). Since B_j has length $O(r)$ and since σ_j is periodic with period $O(r)$, again any infinite path which exits in B_j is contained in an $O(r)$ neighborhood of ∂C along B_j .

Finally, we can glue together the tilings $\sigma_1, \dots, \sigma_m$ by going back and deleting any infinite path in (σ_j, τ_1) which connects to one which would have been deleted in (σ_i, τ_1) for all other $i \neq j$. Since the number of tubes m is constant, in the end we have a tiling τ of C where we have deleted infinite paths of (τ, τ_1) in a neighborhood of width at most constant-order in n (concretely $1000 + O(r)$, where r is constant in n) in distance measuring in edges of $\frac{1}{n}\mathbb{Z}^3$, corresponding to a neighborhood of Euclidean width of $O(r/n)$.

Since all channels C are tiled so that they have only $-\eta_1$ tiles in a neighborhood of ∂C , we can put them together. Therefore we have constructed a tiling $\tau \in T_n(\tilde{R})$.

Bounding the final distance. To emphasize the dependence on n , let τ^n be the tiling at scale n constructed above and let τ_1^n be the $-\eta_1$ brickwork pattern at scale n . On one hand, the total flow of $\tilde{g} + \eta_1$ over any $X_j \in \mathcal{X}$ is

$$\text{vol}(X_j)(v_j + \eta_1).$$

We claim that the double dimer flow $f_{(\tau^n, \tau_1^n)} = f_{\tau^n} - f_{\tau_1^n}$ has the same total flow, up to an $O(n^{-1})$ error. To explain the order of error, recall that for a scale n tiling flow, each

$1/n^3$ -volume lattice cube has flow of order $1/n^3$. The error comes from the region around the boundary of the channels where some infinite paths were deleted and replaced with $-\eta_1$ tiles. The number of lattice sites on the boundary of the channel is order n^2 , and the region has width constant order in n in lattice cubes from $\frac{1}{n}\mathbb{Z}^3$, so the amount of deleted flow in this region has order $n^2/n^3 = 1/n$. Therefore there is a constant K such that

$$\left| \text{vol}(X_j)(\tilde{g}_j + \eta_1) - \sum_{e \in \frac{1}{n}\mathbb{Z}^3, e \cap X_j \neq \emptyset} (f_{\tau^n} - f_{\tau_1^n})(e) \right| < Kn^{-1}.$$

There is also a constant $C = C(\tilde{R})$ such that $M = C\epsilon^{-3}$ (recall that $M = |\mathcal{X}|$ is the number of tetrahedra in the mesh). By Lemma 5.2.7 applied to the partition X_1, \dots, X_M of \tilde{R} ,

$$d_W(f_{\tau^n} - f_{\tau_1^n}, \tilde{g} + \eta_1) < 3M(10\epsilon^4 + Kn^{-1}) < 30C\epsilon + 3CK\epsilon^{-3}n^{-1}.$$

Taking n large enough so that $1/n < \epsilon^4$, this becomes a bound which is a constant times ϵ . A few applications of the triangle inequality and the ‘‘mass shift’’ property of Wasserstein distance, i.e. that $\mathbb{W}_1^{1,1}(\mu, \nu) = \mathbb{W}_1^{1,1}(\mu + \rho, \nu + \rho)$ (see Lemma 5.1.1), give that

$$d_W(f_{\tau^n}, \tilde{g}) < d_W(f_{\tau^n} - f_{\tau_1^n}, \tilde{g} + \eta_1) + d_W(f_{\tau_1^n}, -\eta_1).$$

Since $d_W(f_{\tau_1^n}, -\eta_1) \rightarrow 0$ as $n \rightarrow \infty$, we can make this as small as needed as $n \rightarrow \infty$. Therefore we can choose ϵ small enough and n large enough given δ so that $\tau^n \in T_n(\tilde{R})$ has $d_W(f_{\tau^n}, \tilde{g}) < \delta/2$. Restricting τ^n to R completes the proof. \square

8.5 Soft boundary lower bound

With the machinery developed in the previous section we can now prove the *soft boundary lower bound*, namely Theorem 8.1.10. In particular we show that for (R, b) with b extendable outside, then for any $g \in AF(R, b)$,

$$\lim_{\delta \rightarrow 0} \liminf_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \geq \text{Ent}(g).$$

Recall that μ_n is counting measure on $TF_n(R, b, \theta_n)$, the set of free boundary tiling flows on R at scale n with boundary values within θ_n of b . The main idea of the proof is to show that from the one free boundary tiling flow $f_\tau \in A_\delta(g) \cap TF_n(R)$ constructed in previous section (Theorem 8.4.1), we can use the patching theorem (Theorem 6.3.5) to show that there are actually many tiling flows in $A_\delta(g)$.

Proof of Theorem 8.1.10. By Proposition 8.3.1, there exists $\delta_1 > 0$ such that there is a δ_1 -tetrahedral mesh $\mathcal{X} = \{X_1, \dots, X_M\}$, region $\tilde{R} = \cup_{X \in \mathcal{X}} X$ containing R , and an asymptotic flow $\tilde{g} \in AF(\tilde{R})$ taking constant values on tetrahedra in \mathcal{X} with $d_W(g, \tilde{g}) < \delta/2$ so that

$$A_{\delta/2}(\tilde{g}) \subset A_\delta(g).$$

Let $\tilde{g}_i := \tilde{g}|_{X_i}$. Computing directly,

$$\text{Ent}(\tilde{g}) = \frac{1}{\text{Vol}(R)} \sum_{i=1}^M \text{Vol}(X_i) \text{ent}(\tilde{g}_i).$$

On the other hand by Proposition 7.6.3,

$$\text{Ent}(\tilde{g}) = \text{Ent}(g) + o_\delta(1).$$

Using the shining light construction from the proof of Theorem 8.4.1, for any n large enough there exists a tiling $\tau \in T_n(R)$ such that $f_\tau \in A_{\delta/2}(\tilde{g})$ has a particular form. Let C_1, \dots, C_K denote the channels in the shining light construction. For each tetrahedron X and channel C that intersect, $X \cap C$ is a tube. As in the proof of Theorem 8.4.1, $\tau|_{X \cap C}$ is periodic at a scale independent of n , and has mean current $\tilde{g}_X \in \text{Int}(\mathcal{O})$ on $X \cap C$ outside a neighborhood of $\partial(X \cap C)$ of width constant order in n .

We choose $\epsilon \ll \delta_1$, and partition the interior of $X \cap C$ (where τ has mean current in $\text{Int}(\mathcal{O})$) into small cubes with side length $\leq \epsilon$. For each $i = \{1, \dots, M\}$, call the pieces of the partition contained in X_i

$$\{Q_1^i, \dots, Q_{k_i}^i\}_{i=1}^M.$$

For any (i, k) pair, $\tau|_{Q_k^i}$ is periodic. Recall that Q_k^i has diameter $< \epsilon$. For $\epsilon_1 \ll \epsilon$, for n large enough $\tau|_{\partial Q_k^i}$ is ϵ_1 -nearly-constant with value \tilde{g}_i (see Definition 6.3.3), so it satisfies the conditions for the outer boundary condition in the patching theorem (Theorem 6.3.5). Fix $c \in (0, 1)$. For each (i, k) pair, we choose an EGM $\mu_{i,k}$ of mean current \tilde{g}_i (these exist by Corollary 7.5.4). Since $\tilde{g}_i \in \text{Int}(\mathcal{O})$, a sample from $\mu_{i,k}$ satisfies the conditions of Theorem 6.3.5 with probability going to 1 as $n \rightarrow \infty$ (Corollary 6.6.5). Therefore by Theorem 6.3.5, for n large enough, with probability $(1 - c)$, τ restricted to ∂Q_k^i can be patched with a sample σ from $\mu_{i,k}$ on an annulus of width cn . By the ergodic theorem, given any $\epsilon_2 > 0$, for n large enough we can assume that

$$d_W(f_\sigma|_{Q_k^i}, \tilde{g}_i|_{Q_k^i}) < \epsilon_2 \quad (64)$$

with probability $1 - c$. Therefore with probability $1 - 2c$, Equation (64) holds and σ can be patched with τ .

Let $\pi_{i,k,n}$ denote uniform measure on tilings σ of Q_k^i at scale n with $\sigma|_{\partial Q_k^i} = \tau$ and satisfying Equation (64), and let $Z_n(Q_k^i)$ be its partition function. The additional constraint that Equation (64) is satisfied does not change the exponential order of the number of tilings, hence by Proposition 6.8.3, we get the following consequences for entropy:

$$(1 + O(c))h(\mu_{i,k}) \leq n^{-3}\text{Vol}(Q_k^i)^{-1}H(\pi_{i,k,n}) = n^{-3}\text{Vol}(Q_k^i)^{-1} \log Z_n(Q_k^i).$$

By Lemma 5.2.7 applied to the partition $\{Q_1^i, \dots, Q_{k_i}^i\}_{i=1}^M \cup \{R \setminus \cup_{i=1}^M \cup_{k=1}^{k_i} Q_k^i\}$, if $\sigma \in T_n(R)$ is a free boundary tiling of R whose restrictions to each Q_k^i are in the support of $\pi_{i,k,n}$ then using Equation (64),

$$d_W(f_\tau, f_\sigma) < 3\epsilon^{-3}(10\epsilon^4 + \epsilon_2) + 3C(\partial R)\epsilon$$

where $C(\partial R)$ is a constant depending only on R . In particular, choosing $\epsilon_2 = \epsilon^4$ and taking ϵ sufficiently small, $f_\sigma \in A_{\delta/2}(\tilde{g})$. By Theorem 5.5.7 (uniform continuity of the boundary value operator $T(\cdot, \partial R)$), we can choose ϵ small enough so that the boundary values $T(f_\sigma, \partial R)$ are within θ_n of b for all σ in the support of $\pi_{i,k,n}$. Therefore for n large enough,

$$\mu_n(A_{\delta/2}(\tilde{g})) \geq \prod_{i=1}^M \prod_{k=1}^{k_i} Z_n(Q_k^i) \geq \prod_{i=1}^M \prod_{k=1}^{k_i} \exp\left(n^3 \text{Vol}(Q_k^i) h(\mu_{i,k})(1 + O(c))\right).$$

Recall that $v_n^{-1} = n^{-3}\text{Vol}(R)^{-1}$. Since $\mu_{i,k}$ is an EGM of mean current $\tilde{g}_i \in \text{Int}(\mathcal{O})$, $h(\mu_{i,k}) = \text{ent}(\tilde{g}_i)$ by Theorem 7.5.2. Rearranging and taking into account the $O(\epsilon)$ proportion of each

tetrahedron $X \in \mathcal{X}$ that is not included in the patched regions,

$$\begin{aligned} v_n^{-1} \log \mu_n(A_{\delta/2}(\tilde{g})) &\geq \sum_{i=1}^M \sum_{k=1}^{k_i} \frac{\text{Vol}(Q_k^i)}{\text{Vol}(R)} \text{ent}(\tilde{g}_i)(1 + O(c)) \\ &\geq \sum_{i=1}^M \frac{\text{Vol}(X_i)}{\text{Vol}(R)} \text{ent}(\tilde{g}_i)(1 + O(c))(1 - O(\epsilon)) \\ &= \text{Ent}(\tilde{g})(1 + O(c))(1 - O(\epsilon)). \end{aligned}$$

Recall that ϵ is the size of the patched regions, and $c < \epsilon$ is the patching error. All of ϵ, c, ϵ_2 are much smaller than $\delta > 0$, and go to 0 as $\delta \rightarrow 0$. In particular for any fixed $\delta > 0$, and $\epsilon > 0$ fixed small enough given $\delta > 0$, there is a $n(\delta)$ such that for all $n > n(\delta)$,

$$\liminf_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \geq \liminf_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_{\delta/2}(\tilde{g})) \geq \text{Ent}(\tilde{g}) + O(\epsilon) = \text{Ent}(g) + o_\delta(1).$$

Taking $\delta \rightarrow 0$ completes the proof. \square

8.6 Generalized patching and hard boundary lower bound

To prove the hard boundary lower bound (Theorem 8.2.8), we need one more tool. Recall that $\bar{\rho}_n$ is the uniform probability measure on tilings of a fixed region $R_n \subset \frac{1}{n}\mathbb{Z}^3$.

The shining light construction (Theorem 8.4.1) shows that for any $\delta > 0$ and $g \in AF(R, b)$, for n large enough there exists a free-boundary tiling $\tau \in T_n(R)$ such that $d_W(f_\tau, g) < \delta$. For hard boundary conditions, we need to know that every $g \in AF(R, b)$ can be approximated by a tiling of the *fixed* region R_n . To do this, we prove a *generalized patching theorem* (Theorem 8.6.2).

Let $B_n = [-n, n]^3$. Recall that the patching theorem (Theorem 6.3.5) says that if two tilings τ_1, τ_2 are *nearly constant* with value $s \in \text{Int}(\mathcal{O})$ (Definition 6.3.3), then for any $c > 0$ there is n large enough that we can patch together $\tau_2|_{B_{(1-c)n}}$ and $\tau_1|_{\mathbb{Z}^3 \setminus B_n}$ by tiling the width- cn annulus between them.

In Section 6 where the regular patching theorem (Theorem 6.3.5) was proved, all our tools were combinatorial, and we thought of tilings τ of \mathbb{Z}^3 without rescaling to $\frac{1}{n}\mathbb{Z}^3$. Here we look at the tileability of more general “annular regions,” where the tilings are rescaled to live in $\frac{1}{n}\mathbb{Z}^3$.

Let $R \subset \mathbb{R}^3$ be a compact set which is the closure of a connected domain and has piecewise-smooth boundary ∂R (i.e., the sort of region to which our LDP applies). For any small $c > 0$, we define

$$R^c = \{x \in R : d(x, \partial R) \geq c\}.$$

The set $R \setminus R^c$ is an annular region. On the discrete side, given a free-boundary tiling $\tau \in T_n(R)$, it restricts to $\tau' \in T_n(R^c)$ which is a free boundary tiling of R^c . We let $R_n^c \subset \frac{1}{n}\mathbb{Z}^3$ be the region covered by τ' . Given another free-boundary tiling $\sigma \in T_n(R)$, let $R_n \subset \frac{1}{n}\mathbb{Z}^3$ be the region covered by σ . The region $A = R_n \setminus R_n^c \subset \frac{1}{n}\mathbb{Z}^3$ is the type of annular region we study here.

Let (R, b) be flexible and suppose that $R_n \subset \frac{1}{n}\mathbb{Z}^3$ is a sequence of regions satisfying the conditions of the hard boundary LDP (Theorem 8.2.4), i.e. regions with tileable boundary

values b_n on ∂R converging to b in $\mathbb{W}_1^{1,1}$. To prove any $g \in AF(R, b)$ can be approximated by f_τ with τ a tiling of R_n , we show that we can *patch together* suitable tilings on annuli of the form $R_n \setminus R_n^c$, with hard boundary condition on the outside.

Definition 8.6.1. We say that a flow $g \in AF(R, b)$ is *flexible* if g satisfies the condition that for any compact set $D \subset \text{Int}(R)$, $\overline{g(D)} \subset \text{Int}(\mathcal{O})$.

The pair (R, b) is *flexible* (see Definition 8.2.1 and Lemma 8.2.2) if and only if there exists $g \in AF(R, b)$ which is flexible.

Theorem 8.6.2 (Generalized patching theorem). *Fix $c > 0$. Let (R, b) be flexible, with b a boundary asymptotic flow which is extendable outside. Let $R_n \subset \frac{1}{n}\mathbb{Z}^3$ be a sequence of regions with tileable boundary values b_n on ∂R converging to b in $\mathbb{W}_1^{1,1}$. Let σ_n be a sequence of tilings of R_n .*

Let $\tau_n \in T_n(R)$ be a sequence of tilings such that $d_W(g, f_{\tau_n}) \rightarrow 0$ as $n \rightarrow \infty$ for $g \in AF(R, b)$ flexible. Let τ'_n be τ_n restricted to a free boundary tiling of R^c , and let $R_n^c \subset \frac{1}{n}\mathbb{Z}^3$ be the cubes covered by τ'_n .

For n large enough, $R_n \setminus R_n^c$ is tileable.

Remark 8.6.3. The *flexible* condition here is analogous to the $s \in \text{Int}(\mathcal{O})$ condition in the original patching theorem (Theorem 6.3.5). The generalized patching theorem is the reason the hard boundary LDP requires that (R, b) is flexible.

Before we prove this, we explain how it can be used to prove the hard boundary lower bound (Theorem 8.2.8). First, from the generalized patching theorem, it is straightforward to prove the fixed boundary analog of Theorem 8.4.1.

Corollary 8.6.4. *Suppose that (R, b) is flexible and b is a boundary asymptotic flow which is extendable outside. Fix a sequence of regions R_n with tileable boundary values b_n on ∂R converging to b in $\mathbb{W}_1^{1,1}$. For any $\delta > 0$ and any $g \in AF(R, b)$, there is n large enough such that there exists a tiling τ of R_n with $d_W(f_\tau, g) < \delta$.*

Proof. Since (R, b) is flexible, there exists $g_0 \in AF(R, b)$ which is flexible. For any $\epsilon > 0$, the new flow $g_\epsilon = \epsilon g_0 + (1 - \epsilon)g$ satisfies $d_W(g, g_\epsilon) < C\epsilon$ for some constant $C > 0$. Taking ϵ small enough, we can guarantee that $d_W(g, g_\epsilon) < \delta/2$.

Since g_ϵ is flexible and R_n is tileable, by Theorem 8.6.2 for n large enough there exists a tiling τ of R_n such that $d_W(g_\epsilon, f_\tau) < \delta/2$. By the triangle inequality $d_W(g, f_\tau) < \delta$, which completes the proof. \square

Adding Corollary 8.6.4 as the first step, the proof of Theorem 8.2.8 (hard boundary lower bound) is the same as the proof of Theorem 8.1.10 (soft boundary lower bound).

Proof of Theorem 8.2.8. By Corollary 8.6.4, given any $\delta > 0$, for n large enough and any $g \in AF(R, b)$ we can find a tiling τ of R_n such that $d_W(f_\tau, g) < \delta$. Further, this tiling is of the form given in the shining light construction (proof of Theorem 8.4.1), other than in an annulus of width $c \in (0, 1)$ where c can be taken arbitrarily small. The remainder of the proof is identical to the proof of Theorem 8.1.10, where we use the regular patching theorem to patch in samples from ergodic Gibbs measures of appropriate mean currents. \square




Figure 37: 2D schematic picture for the proof of the generalized patching theorem.

It remains to prove the generalized patching theorem (Theorem 8.6.2). The proof is structurally analogous to the proof of the regular patching theorem for cubes $B_n = [-n, n]^3$ (Theorem 6.3.5), and relies on a sequence of lemmas. We give an outline of the main ideas to explain where each of the lemmas is used, accompanied by the schematic picture in Figure 37. We then state and prove each of the lemmas, followed by a proof of Theorem 8.6.2.

Note that there are some superficial changes between the results here and their analogs in Section 6, since here our regions $R_n \subset \frac{1}{n}\mathbb{Z}^3$ instead of $B_n \subset \mathbb{Z}^3$. Basically this corresponds to a change in units. We introduce a few new pieces of notation to make it easier to work with the tilings τ of $\frac{1}{n}\mathbb{Z}^3$ instead of \mathbb{Z}^3 .

- If Q is a discrete surface built out of lattice squares in $\frac{1}{n}\mathbb{Z}^3$, we define $\text{area}_n(Q)$ to be the number of lattice squares on Q . This is n^2 times the Euclidean area of Q . If Q is a surface not built of lattice squares, we can still use $\text{area}_n(Q)$ to mean the Euclidean area of Q times n^2 .
- If ℓ is a discrete curve built out of edges of squares in $\frac{1}{n}\mathbb{Z}^3$, we defined $\text{length}_n(\ell)$ to be number of lattice edges in L . This is n times the Euclidean length of ℓ .
- If τ is a tiling of $\frac{1}{n}\mathbb{Z}^3$, then the tiling flow f_τ and pretiling flow v_τ are typically *rescaled* so that for $e \in \frac{1}{n}\mathbb{Z}^3$, $v_\tau(e) = \pm 1/n^3$ or 0, and $f_\tau(e) = \pm 5/6n^3$ or $\pm 1/6n^3$. We define \tilde{v}_τ to be *unrescaled flow* $\tilde{v}_\tau(e) = \pm 1$ or 0 for $e \in \frac{1}{n}\mathbb{Z}^3$, and similarly $\tilde{f}_\tau(e) = \pm 5/6$ or $\pm 1/6$.

The proof of Theorem 8.6.2 uses a mixture of combinatorial results like in Section 6 and more analytic results about Wasserstein distance, which are for rescaled tiling flows. These pieces of notation make it easier to go between these points of view, and to explain the purely combinatorial arguments in a way more analogous to Section 6.

The main combinatorial tool is again Hall's matching theorem (Theorem 6.1.3), which says that if $A = R_n \setminus R_n^c$ is not tileable, then there exists a counterexample set $U \subset A$ which proves it. The *interior boundary* S of U is without loss of generality a minimal monochromatic discrete surface with number of squares from $\frac{1}{n}\mathbb{Z}^3$ on the surface bounded above and below by a constant times n^2 (Lemma 8.6.5). The first step is to *indent* slightly and let $A_{\text{mid}} \subset A$

be a slightly smaller annulus where U is well-behaved (Lemma 8.6.6). We then define $U_{\text{mid}} = U \cap A_{\text{mid}}$.

We find a *test tiling* τ_{SL} using a shining light construction (Definition 8.6.8) and show that it satisfies a flow bound (Lemma 8.6.9). This is where we use the condition that g is flexible. Using Lemma 8.6.9, we show that there is a constant $K \in (0, 1)$ such that $\text{imbalance}(U'_{\text{mid}}) \leq -Kn^2 + O(n)$, where U'_{mid} is U_{mid} plus a few cubes from the rest of U (determined by τ_{SL}). This “slack” corresponds to flow from τ_{SL} which exits through the boundary of U_{mid} in the interior, see the pink arrows in Figure 37.

It remains to bound the imbalance in $U_{\text{shell}} := U \setminus U'_{\text{mid}}$. This we break into two pieces:

- Regions of U_{shell} contained in nice “cylinders” connecting ∂A to ∂A_{mid} . These are the regions where there are black and blue arrows in Figure 37, the sides of the cylinder are the orange regions which we call the “ribbon surface.”
- The rest of U_{shell} , which we call the “leftover region.”

Up to error related to the orange area in Figure 37 (the ribbon surface), we show that the imbalance of U_{shell} is the same as the imbalance in the cylinder regions. Finally we relate the imbalance in the cylinders to the flux of tiling flows f_{σ_n} , f_{τ_n} , and $f_{\tau_{\text{SL}}}$, which we can bound using Wasserstein convergence considerations using Lemma 5.2.8, up to an error proportional to the green area in Figure 37.

We now proceed to the lemmas. We first note that Lemma 6.5.2 for cubes has an analog for general regions. The only difference is that the constants c_1, c_2 change since they can depend on the regions.

Lemma 8.6.5. *Let $R_n \subset \frac{1}{n}\mathbb{Z}^3$ be a region of diameter in lattice squares at least n such that R_n are regions approximating a fixed region $R \subset \mathbb{R}^3$ with ∂R a piecewise smooth surface. Define $A = R_n \setminus R_n^c$ for some $c > 0$. Suppose that $S \subset A$ is a monochromatic minimal discrete surface in $\frac{1}{n}\mathbb{Z}^3$ with connects the inner and outer boundaries of A . Then there exist constant c_1, c_2 independent of S and n such that*

$$c_1 n^2 \geq \text{area}_n(S) \geq c_2 n^2$$

where $c_2 \sim c^2$.

Proof. This proof is completely analogous to the proof of Lemma 6.5.2.

For the upper bound, we use that S is minimal to get a bound which is a constant times the surface area of R_n . Since $R_n \subset \frac{1}{n}\mathbb{Z}^3$ are regions approximating R , and since ∂R is piecewise smooth, $\text{area}(R_n)$ is bounded by a constant times n^2 , where this constant is determined by ∂R .

For the lower bound, we apply Proposition 6.5.1 to a point $p \in S$ which is distance at least $c/3$ from ∂A to get that

$$\text{area}(S) \geq \kappa((c/3)n)^2 = c_2 n^2,$$

where κ is a universal constant coming from the isoperimetric inequality. From this expression we see that c_2 is order c^2 . \square

Lemma 8.6.6 (Generalized indenting lemma). *Let $A = R_n \setminus R_n^c$, and fix $\beta > 0$ small. Let S be a minimal monochromatic surface connecting inner and outer boundaries of A . There exists $\epsilon < c$ independent of S such that the following hold for n large enough:*

1. *Let $A_{a,b}$ denote the annulus between layers ∂R_n^a and ∂R_n^b . Then $\text{area}_n(S \cap A_{\epsilon, 2\epsilon}) < \beta n^2$.*
2. *There exists $a \in (\epsilon, 2\epsilon)$ such that $\text{length}_n(\partial(\partial R_n^a \cap S)) \leq (\beta/\epsilon)n$.*
3. *A “ribbon surface” γ for $S \cap \partial R_n^a$ is a surface connecting ∂R_n^a to ∂R_n with boundary $\partial(S \cap \partial R_n^a)$ on ∂R_n^a . The “ribbon area” is the minimal area_n of a ribbon surface, and is bounded by $2\beta n^2$.*

The analogous bounds hold for some $a' \in (c - 2\epsilon, c - \epsilon)$.

Remark 8.6.7. For two surfaces A, B , the set $A \cap B$ is either a surface (2-dimensional), a curve (1-dimensional), or a combination of the two. In any of these cases, we take $\partial(A \cap B)$ to mean that we take union of the curve part of $A \cap B$ and the boundary of the surface part of $A \cap B$.

Proof. By Lemma 8.6.5, $c_2 n^2 \leq \text{area}_n(S) \leq c_1 n^2$. For a given $b > 0$, divide a band of the form $A_{0,b}$ in half. Each time we divide, by Lemma 8.6.5 both the halves of S have area bounded below by a fixed constant times b^2 . On the other hand, one of the halves of S can have at most $1/2$ the original area. Iterating this, we can find $\epsilon > 0$ small enough so that the outer band after we split has area at most βn^2 .

By the pigeonhole principle, there exists a layer an between ϵn and $2\epsilon n$ where

$$\text{length}_n(\partial(\partial R_n^a \cap S)) \leq (\beta n^2)/(\epsilon n) = (\beta/\epsilon)n.$$

Given this, we can find a ribbon surface γ (not necessarily built from lattice squares) with $\text{area}_n(\gamma) \leq (\beta/\epsilon)n(2\epsilon n) = 2\beta n^2$, and hence the ribbon area is bounded by $2\beta n^2$. \square

Suppose that $g \in AF(R, b)$ is flexible. For any $\delta > 0$, we can find a piecewise-constant flow \tilde{g} with $d_W(g, \tilde{g}) < \delta$ (Proposition 8.3.1). More precisely, \tilde{g} is piecewise-constant on a mesh of small tetrahedra \mathcal{X} . The region $\tilde{R} = \cup_{X \in \mathcal{X}} X$ contains R , and $\tilde{g} \in AF(\tilde{R})$. In the proof of Theorem 8.4.1, to construct a tiling approximation $\tau \in T_n(R)$ of g for n large, we construct a tiling approximation $\tilde{\tau} \in T_n(\tilde{R})$ of \tilde{g} for n large, and restrict it to R .

The flexible condition passes from g to \tilde{g} as follows. For any compact set $D \subset \text{Int}(R)$, g flexible means that $\overline{g(D)} \subset \text{Int}(\mathcal{O})$. In particular there is a constant $k_D \in [0, 1)$ such that if $x \in D$ then

$$|g(x)|_1 \leq k_D < 1,$$

where $|\cdot|_1$ denotes the L^1 norm. When g is flexible, we can choose \tilde{g} so that for any compact set $D \subset \text{Int}(R)$, for all tetrahedra $X \in \mathcal{X}$ such that $X \subset D$, $|\tilde{g}_X|_1 \leq k_D$. When this holds, we say that \tilde{g} is a *piecewise constant approximation of g inheriting the flexible condition*.

For any $\epsilon > 0$ and $D \subset \text{Int}(R)$ compact, for $\delta > 0$ small enough we can find \tilde{g} which is piecewise-constant on a scale δ tetrahedral mesh \mathcal{X} , has $d_W(g, \tilde{g}) < \epsilon$, and inherits the flexible condition on D , i.e. $|\tilde{g}_X| \leq k_D$ for all $X \in \mathcal{X}$ such that $X \subset D$.

In the shining light construction in the proof of Theorem 8.4.1, we cut the tetrahedra into tubes (with thin space between them), and construct the tiling approximation $\tilde{\tau}$ of \tilde{g} by filing

the tubes in X with periodic tilings of mean current approximating \tilde{g}_X on X , and fill the thin area in between with all $-\eta_1$ tiles. There is some maximum period r that we use to construct the periodic tilings on the tubes, and r is independent of n . For all n , at all the sites where $\tilde{\tau}$ does not have mean current \tilde{g} , $\tilde{\tau}$ looks locally like the $-\eta_1$ brickwork pattern. The width of the region containing the places where $\tilde{\tau}$ looks like the $-\eta_1$ brickwork pattern is $O(r)$ and therefore independent of n .

Definition 8.6.8 (Shining light measures). Let $g \in AF(R, b)$ and let \tilde{g} be an approximation as discussed above. For each n , let $\tilde{\tau}_n \in T_n(\tilde{R})$ be a tiling produced by the shining light construction with \tilde{g} , where the periodic tilings in the tubes have maximum period r . Fix a large constant $C = O(r)$. We define a *sequence of shining light measures* λ_n for g using \tilde{g} so that for each n , λ_n is uniform measure on tilings of the form $(\tilde{\tau}_n + x) \upharpoonright_R \in T_n(R)$ for $x \in \mathbb{R}^3$ with $|x| \leq C$.

Using this, we prove a lemma analogous to Lemma 6.6.7 from Section 6. Instead of a result for ergodic measures of mean current $s \in \text{Int}(\mathcal{O})$, this lemma is for a sequence of shining light measures λ_n for a flexible flow g . This lemma is why the *flexible* condition is needed for generalized patching and hence for the hard boundary LDP.

Lemma 8.6.9 (Shining light flow bound). *Let $D \subset \text{Int}(R)$, and let $g \in AF(R, b)$ be flexible. Let \tilde{g} be piecewise constant approximation of g on a tetrahedral mesh \mathcal{X} such that \tilde{g} inherits the flexible condition, and such that the union of tetrahedra $\mathcal{D} \subset \mathcal{X}$ covering D is still contained in $\text{Int}(R)$.*

Let S be a monochromatic black surface in $\frac{1}{n}\mathbb{Z}^3$ with boundary ∂S and λ_n be a sequence of shining light measures as in Definition 8.6.8 for \tilde{g} . Let Θ_D be the collection of odd cubes adjacent to $S \cap D$. Let $N = \text{area}_n(S \cap D)$ be the number of squares from $\frac{1}{n}\mathbb{Z}^3$ on $S \cap D$. Then there is a constant $K_D \in (0, 1)$ independent of S such that for all n large enough,

$$\mathbb{E}_{\lambda_n} [|\text{flux}(\tilde{v}_\tau, S \cap D)|] \geq K_D |\Theta_D| + O(n^{-1}N) + O(\text{length}_n(\partial(S \cap D))).$$

The constant K_D depends only on D and g . In particular, it is independent of \tilde{g} , as long as \tilde{g} inherits the flexible condition and is constructed on a small enough mesh \mathcal{X} .

Remark 8.6.10. Recall that for τ a tiling in $\frac{1}{n}\mathbb{Z}^3$, the flow \tilde{v}_τ is the *non-rescaled* pretiling flow.

Proof. Restricted to any tetrahedron $X \in \mathcal{X}$, λ_n samples a tiling which is periodic with mean current \tilde{g}_X up an $O(n^{-1})$ error. In particular, there are probabilities $p_1(X), \dots, p_6(X)$ such that the probability of seeing a tile of type i in τ restricted to X sampled from λ_n is $p_i(X) + O(n^{-1})$ for $i = 1, \dots, 6$. Let N_1, \dots, N_6 denote the corresponding six types of squares f on S . Let $N_i(X)$ denote the number of each type of square in S restricted to X .

Recall that $\mathcal{D} \subset \mathcal{X}$ is the collection of mesh tetrahedra $X \in \mathcal{X}$ such that $X \cap D \neq \emptyset$. By assumption $\cup_{X \in \mathcal{D}} X \subset \text{Int}(R)$. Since \tilde{g} inherits the flexible condition, there is a constant $k_D \in [0, 1)$ such that $|\tilde{g}_X|_1 \leq k_D$ for all $X \in \mathcal{D}$. This is related to the probabilities since

$$\tilde{g}_X = (p_1(X) - p_2(X), p_3(X) - p_4(X), p_5(X) - p_6(X)) \quad (65)$$

Since S is monochromatic, the flux of \tilde{v}_τ across S is equal to minus the number of tiles in τ which cross S . Thus the expected value of the flux (up to sign) can be split as the sum of indicator functions $\mathbb{1}_f$, where $\mathbb{1}_f(\tau)$ is 1 if f is crossed by a tile in τ and 0 otherwise. Then

$$\mathbb{E}_{\lambda_n} [|\text{flux}(\tilde{v}_\tau, S \cap D)|] \geq \sum_{X \in \mathcal{D}} \sum_{f \in S \cap X \cap D} \mathbb{E}_{\lambda_n} [\mathbb{1}_f(\tau)] = \sum_{X \in \mathcal{D}} \sum_{i=1}^6 N_i(X) (p_i(X) + O(n^{-1})).$$

The total number of squares on the surface is $N = \sum_{X \in \mathcal{D}} \sum_{i=1}^6 N_i(X) = \text{area}_n(S \cap D)$. By Lemma 6.6.6, $N_1 + N_2, N_3 + N_4, N_5 + N_6$ are all equal to $N/3$ up to an error of $O(\text{length}_n(\partial(S \cap D)))$. Let

$$p_D = \min_{X \in \mathcal{D}} \max\{\min\{p_1(X), p_2(X)\}, \min\{p_3(X), p_4(X)\}, \min\{p_5(X), p_6(X)\}\}.$$

Since $|\tilde{g}_X|_1 \leq k_D < 1$, for all $X \in \mathcal{D}$ at least four of $p_1(X), \dots, p_6(X)$ must be nonzero, including one from each pair. Combined with Equation (65), one can easily check from this that $p_D \geq (1 - k_D)/6 > 0$. On the other hand by the same arguments as in Lemma 6.6.7,

$$\mathbb{E}_{\lambda_n}[|\text{flux}(\tilde{v}_\tau, S \cap D)|] \geq p_D N/3 + O(n^{-1}N) + O(\text{length}_n(\partial(S \cap D))).$$

Since $|\Theta_D| \leq \text{area}_n(S \cap D) = N$, this completes the proof. The constant $K_D = p_D/3 = (1 - k_D)/18$. As k_D is determined by just g and D , we note that this is independent of the choice of \tilde{g} , as long as \tilde{g} inherits the flexible condition and is constructed on a small enough mesh \mathcal{X} . \square

Equipped with these lemmas, we can now give the proof of the generalized patching theorem.

Proof of Theorem 8.6.2. If $A = R_n \setminus R_n^c$ is not tileable, then by Theorem 6.1.3 there exists a counterexample region $U \subset A$ which has only odd cubes along its interior boundary S , but has

$$\text{imbalance}(U) = \text{even}(U) - \text{odd}(U) > 0.$$

By Corollary 6.2.6, we can assume that the interior boundary $S \subset \partial U$ is a minimal monochromatic discrete surface in $\frac{1}{n}\mathbb{Z}^3$.

By Lemma 8.6.6, for any $\beta > 0$ we can find ϵ and inner and outer layers $a_+ \in (\epsilon, 2\epsilon)$ and $a_- \in (c - 2\epsilon, c - \epsilon)$ such that for $a = a_+$ or $a = a_-$,

$$\text{length}_n(\partial(S \cap \partial R_n^a)) \leq (\beta/\epsilon)n \tag{66}$$

and further such that there is a ribbon surface γ for $\partial(S \cap \partial R_n^a)$ such that

$$\text{area}_n(\gamma) \leq 2\beta n^2. \tag{67}$$

We define the middle annulus $A_{\text{mid}} = R_n^{a_+} \setminus R_n^{a_-}$. Let $U_{\text{mid}} = U \cap A_{\text{mid}}$.

We now fix a compact set $D \subset \text{Int}(R)$. We can assume that D is contained in A_{mid} . We can find a piecewise-constant approximation \tilde{g} of g on a tetrahedral mesh \mathcal{X} which inherits the flexible condition. We let $\mathcal{D} \subset \mathcal{X}$ be the collection of tetrahedra X such that $X \cap D \neq \emptyset$. We can assume that the mesh scale of \mathcal{X} is small enough so that the union of all $X \in \mathcal{D}$ is still contained in $\text{Int}(R)$.

By analogous pigeonhole principle arguments in the indenting lemmas, we can assume that D has width in squares from $\frac{1}{n}\mathbb{Z}^3$ of at least $cn/2$ and has $\text{length}_n(\partial(S \cap D)) = O(n)$.

Let $N = \text{area}_n(S \cap D)$. By Lemma 8.6.9, we can find a sequence of shining light measures λ_n satisfying for n large enough,

$$\mathbb{E}_{\lambda_n}[|\text{flux}(\tilde{v}_\tau, S \cap D)|] \geq K_D |\Theta_D| + O(Nn^{-1}) + O(\text{length}_n(\partial(S \cap D))).$$

By Lemma 8.6.5, $|\Theta_D| \geq N = \text{area}_n(S \cap D) \geq (c_2/4)n^2$. Therefore

$$\mathbb{E}_{\lambda_n}[|\text{flux}(\tilde{v}_\tau, S \cap D)|] \geq \frac{K_D c_2}{4} n^2 + O(n).$$

In particular, we can sample a tiling τ_{SL} from λ_n such that

$$|\text{flux}(\tilde{v}_{\tau_{\text{SL}}}, S \cap D)| \geq \frac{K_D c_2}{4} n^2 + O(n). \quad (68)$$

Let $U_{\tau_{\text{SL}}}$ be the cubes covered by τ_{SL} restricted to U_{mid} . This is a tileable region, so

$$\text{imbalance}(U_{\tau_{\text{SL}}}) = 0.$$

Let $U'_{\text{mid}} \subset U$ be $U_{\tau_{\text{SL}}}$ minus even cubes in $A \setminus U$ with a face on S which are connected to U by τ_{SL} . By Equation (68),

$$\text{imbalance}(U'_{\text{mid}}) \leq -\frac{K_D c_2}{4} n^2 + O(n). \quad (69)$$

Let $U_{\text{shell}} = U \setminus U'_{\text{mid}}$. It remains to bound the imbalance in U_{shell} .

Consider the set $\alpha = U \cap \partial A_{\text{mid}}$. For each connected component α_i of α , we form a closed surface using a cylinder ribbon surface component γ_i of γ and corresponding patch $\alpha'_i \subset \partial A$. Let α' be the union of the α'_i components.

Let V_i be the region enclosed by α_i , γ_i , and α'_i , and let $V = \cup_i V_i$.

The regions $U_{\text{shell}} \setminus V$ are the *leftover regions*. Let W be a connected component of the leftover region. By construction, ∂W either intersects at most one of ∂R_n or ∂R_n^c . Thus the boundary condition on ∂W comes from only one tiling, either τ_n or σ_n .

Suppose it comes from σ_n , i.e. that $\partial W \cap \partial R_n \neq \emptyset$ (the version where it comes from τ_n is identical, we just make a choice for concreteness). Since σ_n can be extended to a tiling of all of R_n , we can extend σ_n to a tiling covering W . Let W_σ be the region covered by the tiles from σ_n which intersect W . Clearly $\text{imbalance}(W_\sigma) = 0$, and $W_\sigma \cap \partial R_n = W \cap \partial R_n$. If a tile in W_σ crosses ∂W , then either

- It crosses $\partial W \cap S$, in which case W_σ contains an even cube which is not contained in $W \subset U$.
- It crosses $\partial W \cap \gamma$ (recall that γ is the ribbon surface). In this case W_σ could have an odd cube which is not in W . However the number of these added cubes over all components W of $U_{\text{shell}} \setminus V$ is bounded by $\text{area}_n(\gamma) \leq 4\beta n^2$.

Therefore

$$\text{imbalance}(U_{\text{shell}} \setminus V) \leq 4\beta n^2.$$

We now show that

$$\text{imbalance}(U_{\text{shell}} \cap V) \leq \text{imbalance}(V) + 4\beta n^2.$$

The $4\beta n^2$ term again comes from the ribbon area. We use ideas analogous to those above. Let Y be a component of $V \setminus (U_{\text{shell}} \cap V)$. First note that Y intersects at most one of ∂R_n and ∂R_n^c , so its boundary condition comes from only one tiling.

Assume σ_n is the tiling which defines the boundary condition on Y , and extend it to a tiling which covers Y . Let Y_σ be the region covered by σ_n tiles which intersect Y . Clearly $\text{imbalance}(Y_\sigma) = 0$. If a tile in Y_σ crosses ∂Y , then it is in one of two cases:

- It crosses $\partial Y \cap S$. Since $Y \subset A \setminus U$, in this case Y_σ contains an odd cube which is not in Y . This makes the imbalance of Y larger.
- It crosses $\partial Y \cap \gamma$ (recall γ is the ribbon surface). In this case Y_σ could have an even cube which is not in Y . However the number of these added cubes over all components Y of $V \setminus (U_{\text{shell}} \cap V)$ is bounded by $\text{area}_n(\gamma) \leq 4\beta n^2$.

Therefore in summary,

$$\text{imbalance}(U_{\text{shell}}) \leq 8\beta n^2 + \text{imbalance}(V).$$

We now relate imbalance to flux. As the proof of Proposition 6.2.3 (where we relate black and white surface area to imbalance), given a set V , we apply the divergence theorem to the reference flow $\tilde{r}(e) = 1/6$ for all e in $\frac{1}{n}\mathbb{Z}^3$ oriented even to odd to get that

$$\text{imbalance}(V) = \text{flux}(\tilde{r}, \partial V) = \frac{1}{6} \left(\text{white}(\partial V) - \text{black}(\partial V) \right).$$

We apply this to the set $V = \cup_i V_i$. By Equation (67), the γ_i contribute at most $4\beta n^2$. Therefore

$$\text{imbalance}(U_{\text{shell}}) \leq \sum_i \text{flux}(\tilde{r}, \partial V_i) \leq 12\beta n^2 + |\text{flux}(\tilde{r}, \alpha) - \text{flux}(\tilde{r}, \alpha')|. \quad (70)$$

For the second inequality, we orient α, α' to always both have inward-pointing normal vector (i.e. inward on ∂A and inward on ∂A_{mid}), meaning one has the opposite normal vector as when we compute flux for ∂V_i . This is why we get a minus sign.

The non-rescaled flow \tilde{f}_τ is the divergence-free version of the pretiling flow \tilde{v}_τ ; related by the equation $\tilde{f}_\tau(e) = \tilde{v}_\tau(e) - \tilde{r}(e)$ for all edges e oriented even to odd. The rescaled version has $f_\tau = \frac{1}{n^3} \tilde{f}_\tau$.

Since the boundary condition on α is given by τ_{SL} and the boundary conditions on α' are given by τ_n on the inner boundary and σ_n on the outer boundary, none of the tiles from the corresponding tilings cross α, α' and hence

$$\text{flux}(\tilde{v}_{\tau_{\text{SL}}}, \alpha) = \text{flux}(\tilde{v}_*, \alpha') = 0,$$

where \tilde{v}_* is equal to \tilde{v}_{τ_n} on the inner boundary of ∂A and is equal to \tilde{v}_{σ_n} on the outer boundary of ∂A . Therefore

$$\text{imbalance}(U_{\text{shell}}) \leq 12\beta n^2 + |\text{flux}(\tilde{f}_{\tau_{\text{SL}}}, \alpha) - \text{flux}(\tilde{f}_*, \alpha')|, \quad (71)$$

where $\tilde{f}_* = \tilde{f}_{\tau_n}$ on the inner boundary of ∂A and $\tilde{f}_* = \tilde{f}_{\sigma_n}$ on the outer boundary of ∂A .

It remains to bound these flux differences, and this is where we use information about the boundary conditions. First note that for any surface X and any tiling τ of $\frac{1}{n}\mathbb{Z}^3$, the flux of the non-rescaled \tilde{f}_τ and the rescaled f_τ are related by:

$$\text{flux}(\tilde{f}_\tau, X) = n^2 \text{flux}(f_\tau, X). \quad (72)$$

We have that the *rescaled versions* of the tiling flows f_{τ_n} and $f_{\tau_{\text{SL}}}$ (rescaled, so without the tildes) converge as $n \rightarrow \infty$ to the $g \in AF(R, b)$ given in the theorem statement, that is,

$$\lim_{n \rightarrow \infty} d_W(f_{\tau_n}, g) = \lim_{n \rightarrow \infty} d_W(f_{\tau_{\text{SL}}}, g) = 0.$$

Recall that $T(\cdot, X)$ denotes the trace operator which takes a flow to its restriction to a surface X . By Theorem 5.5.7, for X fixed and any $f_1, f_2 \in AF(R)$, given any $\delta > 0$ there exists δ_1 such that if $d_W(f_1, f_2) < \delta_1$ then $\mathbb{W}_1^{1,1}(T(f_1, X), T(f_2, X)) < \delta$. Recall also that $T(g, \partial R) = b$, and that we are given that $T(f_{\sigma_n}, \partial R) = b_n$ converges to b in $\mathbb{W}_1^{1,1}$. Given these facts, we can choose n large enough to guarantee the following:

- For the outer boundary ∂R ,

$$\mathbb{W}_1^{1,1}(T(f_{\sigma_n}, \partial R), b) < \delta \quad (73)$$

$$\mathbb{W}_1^{1,1}(T(f_{\tau_{\text{SL}}}, \partial R), b) < \delta. \quad (74)$$

- For the inner boundary ∂R^c ,

$$\mathbb{W}_1^{1,1}(T(f_{\tau_n}, \partial R^c), T(g, \partial R^c)) < \delta \quad (75)$$

$$\mathbb{W}_1^{1,1}(T(f_{\tau_{\text{SL}}}, \partial R^c), T(g, \partial R^c)) < \delta. \quad (76)$$

- Finally, let $\partial R_{\text{mid}} = \partial(R^{a+} \setminus R^{a-})$ be the piecewise smooth surface approximated by $\partial A_{\text{mid}} = \partial(R_n^{a+} \setminus R_n^{a-})$. Then

$$\mathbb{W}_1^{1,1}(T(f_{\tau_{\text{SL}}}, \partial R_{\text{mid}}), T(g, \partial R_{\text{mid}})) < \delta. \quad (77)$$

Recall also that boundary value flows correspond to measures, and note that $\text{flux}(f, X) = T(f, X)(X)$ is the total mass of the measure $T(f, X)$ on X . In particular, for a surface X and tiling τ of $\frac{1}{n}\mathbb{Z}^3$ and $B \subset X$,

$$\text{flux}(f_\tau, B) = T(f_\tau, X)(B).$$

Lemma 5.2.8 applied to measures μ, ν supported on a surface X says that if $\mathbb{W}_1^{1,1}(\mu, \nu) < \delta$, then for any $B \subset X$,

$$\mathbb{W}_1^{1,1}(\mu|_B, \nu|_B) \leq \delta + \delta^{1/2}(C(B) + 1),$$

where $\delta^{1/2}C(B)$ is bounded by 2 times the difference of the area of B and the $\delta^{1/2}$ neighborhood of B within X ; equivalently, by the area of the annulus of width $\delta^{1/2}$ with inner boundary ∂B (see Remark 5.2.9).

To use this, we relate α, α' which are contained in the discrete surfaces ∂A_{mid} and ∂A built out of $\frac{1}{n}\mathbb{Z}^3$ lattice squares, to B, B' on the piecewise smooth surfaces ∂R_{mid} and $\partial R \cup \partial R^c$ respectively.

By Equation (66), $\text{length}_n(\partial\alpha) \leq (2\beta/\epsilon)n$. Correspondingly the Euclidean length is bounded as $\text{length}(\partial\alpha) \leq (2\beta/\epsilon)$. Given this, we can cover $\partial\alpha$ with a collection of cubes \mathcal{C} with Euclidean side length 3ϵ , with $|\mathcal{C}| \leq 2\beta/(3\epsilon^2)$. Since the Euclidean width between ∂A and ∂A_{mid} is less than 2ϵ , \mathcal{C} also covers $\partial\alpha' \subset \partial A$.

The Hausdorff distances between ∂A and $\partial R \cup \partial R^c$ and between ∂A_{mid} and ∂R_{mid} are both bounded by $2/n$. There are corresponding sets $B \subset \partial R_{\text{mid}}$ and $B' \subset \partial R \cup \partial R^c$ which differ from α, α' respectively by Hausdorff distance at most $2/n$. Thus for n large enough, \mathcal{C} also covers B, B' . Since $\partial R \cup \partial R^c$ and ∂R_{mid} are piecewise smooth, there is a constant C' such

that the area of either surface restricted to one of the cubes in \mathcal{C} is at most $C'\epsilon^2$. Since $|\mathcal{C}| \leq 2\beta/(3\epsilon^2)$, there is some constant C such that if $\delta^{1/2} \leq \epsilon$, then

$$\delta^{1/2}C(B) \leq C\beta \quad (78)$$

$$\delta^{1/2}C(B') \leq C\beta. \quad (79)$$

We have that $\text{length}(\partial\alpha) \leq 2\beta/\epsilon$, so the number of $\frac{1}{n}\mathbb{Z}^3$ lattice points along $\partial\alpha$ is bounded by a constant times n (the constant here depends on β/ϵ). Since the Hausdorff distance between α and B is bounded by $2/n$, for any tiling τ of $\frac{1}{n}\mathbb{Z}^3$, the flux of f_τ through a surface is proportional to the number of $\frac{1}{n}\mathbb{Z}^3$ lattice points on the surface times $\frac{1}{n}$. Thus for any tiling τ , since f_τ is divergence-free,

$$|\text{flux}(f_\tau, \alpha) - \text{flux}(f_\tau, B)| \leq O(n^{-1}). \quad (80)$$

By Equation (79), since $\text{length}(\partial B') \leq C(B')$, also have that $\text{length}(\partial B') \leq C\beta/\epsilon$. Since the Hausdorff distance between α' and B' is bounded by $2/n$, the number of $\frac{1}{n}\mathbb{Z}^3$ lattice points on a surface between them is also bounded by a constant times n (the constant here depends on β/ϵ). Thus we analogously get that for any tiling τ of $\frac{1}{n}\mathbb{Z}^3$,

$$|\text{flux}(f_\tau, \alpha') - \text{flux}(f_\tau, B')| \leq O(n^{-1}). \quad (81)$$

Therefore by Lemma 5.2.8, for δ such that $\delta^{1/2} < \epsilon$, Equations (73), (75) to relate f_* to g on B' , plus Equation (79) where we determine the constant $C(B')$, and finally Equation (81) to relate f_* on α' to f_* on B' , we get that

$$|\text{flux}(f_*, \alpha') - \text{flux}(g, B')| \leq \delta + \delta^{1/2} + C\beta + O(n^{-1}), \quad (82)$$

where f_* is f_{σ_n} on the outer boundary of ∂A and f_{τ_n} on the inner boundary. Similarly, using Equations (74), (76), and (77) to relate $f_{\tau_{\text{SL}}}$ to g , plus Equations (80), (78) for the constant $C(B)$ and to relate $f_{\tau_{\text{SL}}}$ on α, B , the test tiling τ_{SL} satisfies analogous bounds on both α and α' :

$$|\text{flux}(f_{\tau_{\text{SL}}}, \alpha') - \text{flux}(g, B')| \leq \delta + \delta^{1/2} + C\beta + O(n^{-1}) \quad (83)$$

$$|\text{flux}(f_{\tau_{\text{SL}}}, \alpha) - \text{flux}(g, B)| \leq \delta + \delta^{1/2} + C\beta + O(n^{-1}). \quad (84)$$

Since g is divergence-free and takes values with norm bounded between -1 and 1 , and B, B' differ from α, α' by Hausdorff distance bounded by $2/n$, Equation (67) implies that

$$|\text{flux}(g, B) - \text{flux}(g, B')| \leq 4\beta + O(n^{-1}). \quad (85)$$

Combining Equation (72) with the above,

$$|\text{flux}(\tilde{f}_{\tau_{\text{SL}}}, \alpha) - \text{flux}(\tilde{f}_*, \alpha')| \quad (86)$$

$$\leq \left[|\text{flux}(f_{\tau_{\text{SL}}}, \alpha) - \text{flux}(g, B)| + |\text{flux}(g, B) - \text{flux}(g, B')| + |\text{flux}(g, B') - \text{flux}(f_*, \alpha)| \right] n^2 \quad (87)$$

$$\leq (2\delta + 2\delta^{1/2} + 2C\beta + O(n^{-1}) + 4\beta)n^2. \quad (88)$$

Combining this with Equation (69) and Equation (71), we get that

$$\text{imbalance}(U) = \text{imbalance}(U'_{\text{mid}}) + \text{imbalance}(U_{\text{shell}})$$

$$\leq -\frac{K_{DC_2}}{4}n^2 + (16 + 2C)\beta n^2 + 2\delta n^2 + 2\delta^{1/2}n^2 + O(n).$$

The factor $K_{DC_2}/4$ and the constant C are fixed independent of δ, β . Implicit here is also the parameter ϵ that we indent by.

Taking ϵ small enough, we can make β as small as needed. The parameter $\delta > 0$ is related to the distance between tiling flows and their limits and is required to satisfy $\delta^{1/2} < \epsilon$, but this can be guaranteed for n large enough. Therefore for n large enough, $\text{imbalance}(U)$ will be non-positive, and hence U is not a counterexample. This completes the proof. \square

8.7 Upper bounds

To complete the proof of the large deviation principles, we need prove the upper bounds, namely Theorem 8.1.11 for the soft boundary LDP and Theorem 8.2.9 for the hard boundary LDP. We show that the soft boundary upper bound implies the hard boundary one, and then prove the soft boundary one.

Lemma 8.7.1. *Theorem 8.1.11 implies Theorem 8.2.9.*

Proof. Recall that μ_n is counting measure on $TF_n(R, b, \theta_n)$ for some sequence of thresholds $(\theta_n)_{n \geq 1}$ with $\theta_n \rightarrow 0$ as $n \rightarrow \infty$ sufficiently slowly.

On the other hand $\bar{\mu}_n$ is counting measure on tilings of fixed regions R_n with scale n tileable boundary value b_n such that $b_n \rightarrow b$ as $n \rightarrow \infty$.

We choose the sequence of thresholds θ_n so that there exists N such that if $n \geq N$ then $\mathbb{W}_1^{1,1}(b_n, b) < \theta_n$. In this case, for any $g \in AF(R, b)$, for n large enough,

$$\bar{\mu}_n(A_\delta(g)) \leq \mu_n(A_\delta(g)).$$

Therefore if Theorem 8.1.11 holds, then for all $g \in AF(R, b)$,

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} v_n^{-1} \log \bar{\mu}_n(A_\delta(g)) \leq \lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \leq \text{Ent}(g).$$

This completes the proof of Theorem 8.2.9 from Theorem 8.1.11. \square

It remains to prove Theorem 8.1.11, namely that for any $g \in AF(R, b)$,

$$\lim_{\delta \rightarrow 0} \limsup_{n \rightarrow \infty} v_n^{-1} \log \mu_n(A_\delta(g)) \leq \text{Ent}(g).$$

The main idea is ‘‘coarse graining’’, i.e. that on a very small box, a uniform random tiling of R looks approximately like a random tiling sampled from a $\mathbb{Z}_{\text{even}}^3$ -invariant Gibbs measure of mean current s , where s is the expected mean current on the box.

Proof of Theorem 8.1.11. In this proof, we assume without loss of generality that R is contained in the unit cube $B = [0, 1]^2$ (this is just to avoid complicating the proof with an extra scaling parameter).

Let $\pi_{n,\delta}$ be the uniform probability measure on the set of tilings τ with tiling flow $f_\tau \in A_\delta(g) \cap TF_n(R)$ and satisfying $\mathbb{W}_1^{1,1}(T(f_\tau), b) < \theta_n$. The purpose of this is so that the partition function of $\pi_{n,\delta}$ is $Z_{n,\delta} = \mu_n(A_\delta(g))$.

Tile \mathbb{R}^3 by translated copies of B , each with a translated copy of R inside it. Let $\Lambda_n = \frac{1}{n}\mathbb{Z}^3 \cap B$. We define a $\mathbb{Z}_{\text{even}}^3$ -invariant measure ν_n on tilings of \mathbb{Z}^3 as follows (this measure can sample tilings with some double tiles and some untiled sites). We take an independent sample from $\pi_{n,\delta}$ on each copy of R , and then average over translations by $x \in \text{even}(\Lambda_n)$. The measure ν_n samples tilings that are perfect matchings on the interior of each copy of R . All sites in each copy of $B \setminus R$ which are not covered by a tile connecting it to a copy of R are empty. Two copies of R might intersect on their boundaries (e.g. in the case $R = B$), in which case ν_n can sample double tiles. However the fraction of possible sites where ν_n samples double tiles is bounded by the fraction of sites in ∂B , namely

$$\frac{6n^2}{n^3} = \frac{6}{n}.$$

We define a subsequential limit

$$\nu := \lim_{j \rightarrow \infty} \nu_{n_j}.$$

Note that ν is a $\mathbb{Z}_{\text{even}}^3$ -invariant measure on tilings of \mathbb{Z}^3 , allowed to have untiled sites. It can be written as a weighted average of a measure on dimer tilings and the empty ensemble.

Let $\nu_{n,0}$ be defined analogously to ν_n , but without averaging over translations. Let $\nu_{n,x}$ be the version where all tilings are translated by a fixed $x \in \text{even}(\Lambda_n)$. The Shannon entropy of $\pi_{n,\delta}$ is

$$n^{-3} \text{Vol}(R)^{-1} \log Z_{n,\delta} = v_n^{-1} \log \mu_n(A_\delta(g)).$$

By construction,

$$|\Lambda_n|^{-1} H_{\Lambda_n}(\nu_{n,0}) = n^{-3} \log \mu_n(A_\delta(g)).$$

For any other $x \in \text{even}(\Lambda_n)$, a sample on B contains pieces from up to 8 samples of $\pi_{n,\delta}$. Since ν_n is a uniform measure, and since there are more tilings when we are allowed more double tiles,

$$H_{\Lambda_n}(\nu_{n,x}) \geq H_{\Lambda_n}(\nu_{n,0}) \quad \forall x \in \text{even}(\Lambda_n).$$

The specific entropy of ν_n can be computed using any sequence of boxes Δ_M with $|\Delta_M| \rightarrow \infty$ as $M \rightarrow \infty$. In particular, we can choose $\Delta_M = \Lambda_{Mn}$ so that

$$h(\nu_n) = \lim_{M \rightarrow \infty} |\Lambda_{Mn}|^{-1} H_{\Lambda_{Mn}}(\nu_n).$$

On each of the M^3 copies of Λ_n in Λ_{Mn} , ν_n samples an independent draw from $\pi_{n,\delta}$. Thus

$$h(\nu_n) \geq \lim_{M \rightarrow \infty} M^3 |\Lambda_{Mn}|^{-1} H_{\Lambda_n}(\nu_{n,0}) = |\Lambda_n|^{-1} H_{\Lambda_n}(\nu_{n,0}) = n^{-3} \log \mu_n(A_\delta(g)).$$

On the other hand, since h is upper-semicontinuous,

$$\limsup_{j \rightarrow \infty} h(\nu_{n_j}) \leq h(\nu).$$

Therefore

$$\limsup_{j \rightarrow \infty} v_{n_j}^{-1} \log \mu_{n_j}(A_\delta(g)) \leq \frac{1}{\text{Vol}(R)} h(\nu),$$

and we have reduced the problem to bounding $h(\nu)$. Define φ_n to be the ν_n -expected flow, namely

$$\varphi_n := Z_{n,\delta}^{-1} \sum_{\tau} f_{\tau},$$

where the sum is over tilings τ in the support of $\pi_{n,\delta}$. We define a subsequential limit

$$\varphi := \lim_{j \rightarrow \infty} \varphi_{n_j}.$$

Up to taking additional subsequences we can assume that the subsequences for φ_n and ν_n are the same. Note that $\varphi \in AF(R)$. Since $A_{\delta}(g)$ is convex, Theorem 5.3.1 implies that $\varphi \in \overline{A_{\delta}(g)}$. Therefore

$$\frac{1}{\text{Vol}(R)} \int_R \text{ent}(\varphi(x)) dx = \text{Ent}(\varphi) \leq \sup_{h \in \overline{A_{\delta}(g)}} \text{Ent}(h) = \text{Ent}(g) + o_{\delta}(1).$$

The last equality uses that Ent is upper semi-continuous in the Wasserstein topology (Proposition 7.6.3). This reduces the problem to showing that $\frac{1}{\text{Vol}(R)} h(\nu)$ is bounded by $\text{Ent}(\varphi)$. To this end, we partition B into a collection \mathcal{C} of k^3 smaller cubes of size $1/k^3$. We define a new flow α_k supported in R by, for all $C \in \mathcal{C}$ such that $C \cap R \neq \emptyset$,

$$\alpha_k(x) = \frac{1}{|C \cap R|} \int_{C \cap R} \varphi(y) dy \quad \forall x \in C \cap R.$$

For $x \notin R$, $\alpha_k(x) = 0$. Since ent is concave (Lemma 7.2.1), by Jensen's inequality,

$$\int_{C \cap R} \text{ent}(\alpha_k(x)) dx \geq \int_{C \cap R} \text{ent}(\varphi(x)) dx.$$

On the other hand, α_k converges to φ a.s. and $|\alpha_k| \leq 1$, so α_k converges to φ in L^1 , hence by Corollary 7.6.2,

$$\lim_{k \rightarrow \infty} \text{Ent}(\alpha_k) = \text{Ent}(\varphi).$$

Therefore it is sufficient to show that $\text{Ent}(\alpha_k)$ is an upper bound for all k . We now define $\nu_{n,C}$ to be ν_n but averaged only over the translations $x \in \text{even}(\Lambda_n \cap C)$ (equivalently, conditioned on the origin being in C). For each $C \in \mathcal{C}$, let ν_C be a subsequential limit of $\nu_{n,C}$. The measures ν_C are $\mathbb{Z}_{\text{even}}^3$ -invariant, and we can choose the subsequences so that

$$\nu = k^{-3} \sum_{C \in \mathcal{C}} \nu_C.$$

Therefore

$$h(\nu) = k^{-3} \sum_{C \in \mathcal{C}} h(\nu_C).$$

Note that if $C \cap R = \emptyset$, then ν_C is the empty ensemble and hence in that case $h(\nu_C) = 0$. When $C \cap R \neq \emptyset$, then ν_C splits as the sum of an empty ensemble (corresponding to selecting the origin in $C \setminus C \cap R$) and a measure on dimer tilings (corresponding to selecting the origin in $C \cap R$). In a slight abuse of notation we refer to the mean current of ν_C as the mean current of its component which is a measure on dimer tilings. To bound $h(\nu_C)$ when $C \cap R \neq \emptyset$ we

compute this mean current. Recall from Section 2.2 that $s_0(\tau)$ is the vector of the tile at the origin in τ , and the mean current of a $\mathbb{Z}_{\text{even}}^3$ -invariant measure μ can be computed as

$$s(\mu) = \int_{\Omega} s_0(\tau) d\mu(\tau) = \mathbb{E}_{\mu}[s_0(\tau)].$$

We can also compute the mean current by looking at the expected tile direction over a set of points instead of just looking at the origin. Let $E(\Lambda_n)$ denote the edges in Λ_n oriented from even to odd. We can similarly define

$$s_{n,C}(\tau) := |\text{even}(\Lambda_n \cap C \cap R)|^{-1} \sum_{e \in E(\Lambda_n \cap C \cap R)} f_{\tau}(e)e.$$

Here note that we intersect with R because if the origin is chosen in $C \setminus C \cap R$ we get the empty ensemble, and hence the mean current is not defined. This is the average direction of f_{τ} over $\Lambda_n \cap C \cap R$, and by $\mathbb{Z}_{\text{even}}^3$ -invariance $s(\mu)$ can also be computed

$$s(\mu) = |\text{even}(\Lambda_n \cap C \cap R)|^{-1} \sum_{x \in \text{even}(\Lambda_n \cap C \cap R)} \int_{\Omega} s_0(\tau + x) d\mu(\tau) = \int_{\Omega} s_{n,C}(\tau) d\mu(\tau).$$

Using this, we compute that

$$\mathbb{E}_{\nu_{n,C}}[s_{n,C}(\tau)] = Z_{n,\delta}^{-1} \sum_{\tau} \sum_{e \in E(\frac{1}{n}\mathbb{Z}^3 \cap C \cap R)} |C \cap R|^{-1} f_{\tau}(e)e = \text{avg}_{C \cap R}(\varphi_n).$$

Since ν_C is a subsequential limit of $\nu_{n,C}$ (up to choice of another subsequence),

$$\begin{aligned} s_C := s(\nu_C) &= \int_{\Omega} s_0(\tau) d\nu_C(\tau) = \lim_{j \rightarrow \infty} \int_{\Omega} s_0(\tau) d\nu_{n_j,C}(\tau) = \lim_{j \rightarrow \infty} \int_{\Omega} s_{n_j,C}(\tau) d\nu_{n_j,C}(\tau) \\ &= \lim_{j \rightarrow \infty} \text{avg}_{C \cap R}(\varphi_{n_j}). \end{aligned}$$

On the other hand, since $\varphi_{n_j} \rightarrow \varphi$ in L^1 ,

$$s_C = \limsup_{j \rightarrow \infty} \text{avg}_{C \cap R}(\varphi_{n_j}) = \text{avg}_{C \cap R}(\varphi).$$

Finally we relate $h(\nu_C)$ to $\text{ent}(s_C)$. Recall that $\text{ent}(s) := \max_{\rho \in \mathcal{P}^s} h(\rho)$, where \mathcal{P}^s is the space of $\mathbb{Z}_{\text{even}}^3$ -invariant probability measures on dimer tilings of mean current s . The measure ν_C is a sum of an empty ensemble (corresponding to the origin being chosen in $C \setminus R$) which has zero entropy and a $\mathbb{Z}_{\text{even}}^3$ -invariant on dimer tilings (corresponding to the origin being chosen in $C \cap R$) which has mean current s_C . Thus

$$h(\nu_C) \leq \text{ent}(s_C).$$

Therefore for all $k > 1$,

$$\begin{aligned} \limsup_{j \rightarrow \infty} v_{n_j}^{-1} \mu_{n_j}(A_{\delta}(g)) &\leq \frac{1}{\text{Vol}(R)} h(\nu) = \frac{1}{\text{Vol}(R)} \frac{1}{k^3} \sum_{C \in \mathcal{C}} h(\nu_C) \\ &\leq \frac{1}{\text{Vol}(R)} \frac{1}{k^3} \sum_{C \in \mathcal{C}} \text{Vol}(R \cap C) \text{ent}(s_C) = \text{Ent}(\alpha_k). \end{aligned}$$

Taking $k \rightarrow \infty$, this shows that

$$\limsup_{n_j \rightarrow \infty} v_{n_j}^{-1} \mu_{n_j}(A_{\delta}(g)) \leq \text{Ent}(\varphi) = \text{Ent}(g) + o_{\delta}(1).$$

Since this holds for any convergent subsequence n_j , taking $\delta \rightarrow 0$ completes the proof. \square

9 Open problems

We mentioned in the introduction that there is literature exploring the local move connectivity problem, considering moves such as the “flip” and “trit” illustrated below.



Figure 38: Flip and trit.

Both the flip and the trit amount to finding a cycle in \mathbb{Z}^3 (of length 4 or 6 respectively) that alternates between membership and non-membership in τ , and then swapping the members and non-members. Generally, a *cycle swap* is a swap of an alternating cycle of length k , and a *k-swap* is a cycle swap for which the cycle has length k . It is clear that any two perfect matchings of the same region can be connected by a sequence of such swaps of this form (simply by applying swaps to all of the cycles contained in the union of the two perfect matchings). But it is in general not so clear whether one can get from any matching to any other using only k -swaps for small k .

Problem 9.0.1. Is there a finite K such that for any positive j , m and n (at least one of which is even) it possible to get from any dimer configuration of an $j \times m \times n$ box to any other via sequence of k -swaps with $k \leq K$? Is this possible using only flips and trits?

The examples we have presented in Section 3 already show the answer to both questions is no if one replaces boxes with general simply connected regions, such as those that can be tiled with alternating slabs of brickwork, each oriented a different direction. If we think in terms of the non-intersecting path interpretation from Section 1.5, we can see that the existence of taut patterns like the ones shown there are an obstruction to local move connectedness.

Progress was made on Problem 9.0.1 just after the first draft of this paper was released in [HLT23]. In particular their results show that any dimer tiling of a $j \times m \times n$ box (for jmn even, $j, m, n \geq 2$) admits at least one flip or trit [HLT23, Theorem 1]. See Section 3 for further description of their results.

Problem 9.0.2. What can be said about the convergence rate of the mixing algorithm described in Section 3.3? Is there a more efficient way to sample random perfect matchings of 3D regions?

Problem 9.0.3. Is there a unique ergodic Gibbs measure corresponding to each mean current in the interior of \mathcal{O} ?

Problem 9.0.4. If ν_1 and ν_2 are ergodic Gibbs measures of the same mean current, and (τ_1, τ_2) is sampled uniformly from $(\nu_1 \otimes \nu_2)$, are there necessarily infinitely many infinite paths in the union of τ_1 and τ_2 ?

Problem 9.0.5. What can be said about the *typical fluctuations* of the flow associated to a uniformly random perfect matching of a simple region such as a cube or torus? Do they converge to a natural Gaussian process?

In 2D, Kenyon showed that domino tiling height functions converge in law to the Gaussian free field [Ken00]. This suggests that the discrete gradients of the height functions should converge (at least in some sense) to the gradient of the Gaussian free field. The dual of the

discrete gradient (i.e., the discrete flow) should converge in some sense to the dual of the gradient of the Gaussian field—which can be shown to be equivalent to the field obtained by projecting vector-valued white noise orthogonally onto the space of divergence-free fields. It seems reasonable to conjecture that the same holds in any dimension.

Problem 9.0.6. Is it the case for $d \geq 3$ that the discrete divergence-free flows obtained from uniformly random perfect matchings (on a torus or box, say, or in the \mathbb{Z}^3 Gibbs measure setting) converge in the fine mesh limit to the Gaussian random generalized flow obtained by projecting vector-valued white noise onto the space of divergence-free flows?

Problem 9.0.7. Does there exist a three-dimensional region $R \subset \mathbb{R}^3$ and a boundary condition b for which the Ent maximizer for (R, b) is not unique? We have shown that such a system would have to be “rigid” in the sense defined in the introduction (i.e., there is an interior point x such that for any neighborhood U of x the set $\overline{g(U)}$ must intersect one of the edges of \mathcal{O}). But we have not ruled out the existence of multiple Ent maximizers.

In fact there do exist two dimensional surfaces R where the corresponding Ent maximizer is not unique. Consider the “slanted cylinder” below, where the left and right edges are glued following the numbers in the diagram. Here are two possible tilings of the slanted cylinder.




Figure 39: Two tilings of the slanted cylinder. The left and right edges are glued.

Any tiling of the slanted cylinder consists of a choice of north (N) or east (E) tile for each diagonal, so if the cylinder has height m then it has 2^m distinct tilings. Since there is only one choice to make on each of the diagonal “stripes” (deciding whether to color it blue or orange) the entropy per site tends to zero as the width of the cylinder tends to infinity, and the functions obtained as fine-mesh limits of these constructions are all maximizers of Ent. A slanted cylinder can also be realized as an induced subgraph of \mathbb{Z}^3 as shown below.




Figure 40: Front and back sides of the surface of a cube, with five “stripes” wrapping around it, whose vertices correspond to the squares in Figure 39 and form a slanted cylinder embedded in \mathbb{Z}^3 . Hall’s matching theorem implies that every perfect matching of the set of vertices hit by these stripes is obtained by choosing one of the two possible perfect matchings within each stripe.

If we try to take a fine mesh limit of this example, we get a region with zero volume in \mathbb{R}^3 . The question is whether this kind of phenomenon can arise for regions with non-zero volume that are ordinary subsets of \mathbb{R}^3 (as opposed to, say, 3D analogs of the slanted cylinder). “Thickening” the ribbon-like example above (by taking the union of multiple ribbon layers, taken on different concentric cubes) does not seem to work, as a tiling obtained that way need not be locally frozen (trit moves may be possible at the corners).

Problem 9.0.8. Is there a region $R \subset \mathbb{R}^3$ and boundary asymptotic flow b where the Ent maximizing flow takes values on a face of $\partial\mathcal{O}$ within a strict subset of the interior of R ? Or within a strict subset of all of R ? (Does this happen on the boundary of the Aztec octahedron?) For an example of (R, b) where the limit shape takes values in a face of $\partial\mathcal{O}$ on all of R (i.e., not only a strict subset of the region), see Example 8.2.6.

Problem 9.0.9. Given a region $R \subset \mathbb{R}^3$ and a flow b on ∂R , is there an elegant way to describe the conditions under which $AF(R, b)$ is nonempty? In other words, under what conditions does b admit an extension to R which is an asymptotic flow (measurable, divergence-free, and valued in \mathcal{O})? Recall that if $R \subset \mathbb{Z}^3$ is a discrete region and b is a discrete vector field on $\partial R \subset \mathbb{Z}^3$, then Hall’s matching theorem or the min cut, max flow principle say that b is extendable if and only if there is no *counterexample region* $U \subset R$ such that $S = \partial U \cap R$ is a type of discrete minimal surface, and any extension of b would be required to have too much flow across S . Is there a continuum version of Hall’s matching theorem and the min cut, max flow principle that characterizes when b on ∂R can be extended to an asymptotic flow—i.e., a statement that b is extendable as long as there is no “minimal surface” cut S such that any extension of b would be required to have too much flow across S ? See e.g. [Str10] for discussion of related problems.

A particularly simple case of interest is that where R is a polyhedron and the boundary value b is constant on the faces of the polyhedron.

Problem 9.0.10. Let us try to generalize Aztec prism example from the introduction. Suppose $R \subset \mathbb{R}^3$ is a prism of the form $S \times [0, 1]$ (where S is a two-dimensional region) and b is equal to 0 on the top and bottom faces of the prism. Alternatively, one may identify the top and bottom of the prism, to obtain S cross a circle. We expect that one can show from basic symmetry that the Ent minimizing flow g has zero flow in the vertical direction, that its restriction to a slice $S \times \{x\}$ does not depend on x . Understanding the behavior within this slice is then a two-dimensional flow problem. Is this behavior the same as what one would see for the corresponding two-dimensional dimer model on the slice?

Problem 9.0.11. What can be said about the interfaces between frozen regions on the boundaries of limit shapes (such as those apparent in the figures in the introduction)? How large do the fluctuations tend to be?

Problem 9.0.12. The 2D Aztec diamond has four frozen regions (one for each vertex) and the 3D Aztec octahedron appears to have twelve frozen regions (one for each edge). One might guess that in the k -dimensional analog we would see $4\binom{k}{2}$ frozen regions, one for each co-dimension-two boundary simplex. Can anything along these lines be proved, either in 3D or in higher dimensions?

Problem 9.0.13. In two dimensions, the large deviation theory [CKP01] can be generalized to many other types of random height function models [She05], even though for most of these models we cannot compute ent explicitly. For example, instead of having height differences constrained to $\{3/4, -1/4\}$ as in the 2D dimer model, they could be constrained to some other set, like $\{-1, 1\}$ or $\{-1, 0, 1\}$. That raises a natural question for us. To what other discrete divergence-free flow models in 3D (or in higher dimensions) can the results of this paper be extended? For example, what if instead of restricting the even-to-odd flows to lie in $\{5/6, -1/6\}$ we restrict them to $\{-1, 1\}$ or to some other set? Would the max-flow-min-cut theory available in these settings allow us to complete the steps that relied on Hall’s matching theorem in this paper? Could the “chain swapping” arguments used in this paper be adapted to establish the strict concavity of ent in these settings?

As we mentioned earlier, given a lattice flow v on \mathbb{Z}^3 one can define a discrete “curl” that assigns to each oriented plaquette—which corresponds to an oriented edge of the dual lattice—the flow of v around that plaquette. One can then define a *vector potential function* A_τ on the dual lattice of \mathbb{Z}^3 whose curl corresponds to the flow f_τ on \mathbb{Z}^3 , though A_τ is *a priori* only determined up to the addition of a vector field with curl zero. Restricting the flow f_τ to take values in $\{5/6, -1/6\}$ is then equivalent restricting the curl of A_τ to lie in $\{5/6, -1/6\}$.

Readers familiar with lattice gauge theory (see [Cha19] for a survey) can tell a similar story about a constrained lattice *connection* with gauge group $U(1)$ (the complex unit circle) as follows. Fix some small constant $\alpha \in (0, \pi)$ and constrain the holonomy around every plaquette (oriented clockwise as one looks from the even to the odd incident cube) to lie in $\{e^{5\alpha i/6}, e^{-\alpha i/6}\}$. Then define a domino to be a pair of cubes separated by a plaquette with holonomy $e^{5\alpha i/6}$. Since the product of oriented holonomies around a single cube is zero, each interior cube belongs to exactly one domino, and (up to boundary conditions) one expects a uniformly random constrained connection to correspond to a uniformly random 3D domino tiling.

Problem 9.0.14. Can our large deviation theory be extended to any other types of holonomy-constrained random connections, Abelian or otherwise? Are there other aspects of gauge theory for which this perspective is useful?

Although we have not explained this in detail, we believe that all the arguments of this paper will still apply to the setting where the edges are periodically “weighted” in the manner described in [KOS06]. For example, one might consider a weighting that strongly favors edges whose vertices have the form (x, y, z) and $(x, y, z + 1)$ where z is even. If the weight is strong enough, one can use a standard Peierls argument to show if we are given two independent samples from the minimal-specific-free-energy ergodic Gibbs measure, then there are a.s. no infinite paths in their union.

Problem 9.0.15. If we allow periodic weights, as in [KOS06], what can we say about the phase diagram? Are there some choices of weights for which the double dimer model a.s. contains no infinite paths and others for which it a.s. contains infinitely many infinite path? Are there any other possibilities? Can one say, even on a rough qualitative level, how similar the function ent described here (and its periodically-edge-weighted analogs) will be to the surface tension functions described in [KOS06] (which are interesting algebraic geometry constructions with finitely many singular cusps)? In this generalized setting, can one say anything about the magnitude of the typical fluctuations of a random flow, or how such fluctuations might depend on the edge weights?

References

- [Ada94] C.C. Adams. *The Knot Book*. W.H. Freeman, 1994.
- [Alm86] Fred Almgren. Optimal isoperimetric inequalities. *Indiana University mathematics journal*, 35(3):451–547, 1986.
- [AMS11] Luigi Ambrosio, Edoardo Mainini, and Sylvia Serfaty. Gradient flow of the Chapman-Rubinstein-Schatzman model for signed vortices. *Ann. Inst. H. Poincaré C Anal. Non Linéaire*, 28(2):217–246, 2011.
- [AriBC] Aristotle. On the Heavens, Book III, Part 8, 350 BC.
- [Bed19a] Grigory Bednik. Hopfions in a lattice dimer model. *Physical Review B*, 100(2):024420, 2019.
- [Bed19b] Grigory Bednik. Probing topological properties of a three-dimensional lattice dimer model with neural networks. *Physical Review B*, 100(18):184414, 2019.
- [Bro86] A Z Broder. How hard is it to marry at random? (On the approximation of the permanent). In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, STOC '86, page 50–58, New York, NY, USA, 1986. Association for Computing Machinery.
- [BS94] Robert Burton and Jeffrey E Steif. Non-uniqueness of measures of maximal entropy for subshifts of finite type. *Ergodic Theory and Dynamical Systems*, 14(2):213–235, 1994.
- [CEP96] Henry Cohn, Noam Elkies, and James Propp. Local statistics for random domino tilings of the Aztec diamond. *Duke Mathematical Journal*, 85(1):117 – 166, 1996.
- [Cha19] Sourav Chatterjee. Yang–mills for probabilists. In *Probability and Analysis in Interacting Physical Systems: In Honor of SRS Varadhan, Berlin, August, 2016*, pages 1–16. Springer, 2019.

- [CKP01] Henry Cohn, Richard Kenyon, and James Propp. A variational principle for domino tilings. *Journal of the American Mathematical Society*, 14(2):297–346, 2001.
- [DK95] C Dress and W Krauth. Cluster algorithm for hard spheres and related systems. *Journal of Physics A: Mathematical and General*, 28(23):L597, 1995.
- [DSS10] Daniela De Silva and Ovidiu Savin. Minimizers of convex functionals arising in random surfaces. *Duke Math. J.*, 151(3):487–532, 2010.
- [DZ09] A. Dembo and O. Zeitouni. *Large Deviations Techniques and Applications*. Stochastic Modelling and Applied Probability. Springer Berlin Heidelberg, 2009.
- [FF56] L. R. Ford and D. R. Fulkerson. Maximal flow through a network. *Canadian Journal of Mathematics*, 8:399–404, 1956.
- [FHNQ11] Michael Freedman, Matthew B Hastings, Chetan Nayak, and Xiao-Liang Qi. Weakly coupled non-Abelian anyons in three dimensions. *Physical Review B*, 84(24):245119, 2011.
- [FKMS22] Juliana Freire, Caroline J. Klivans, Pedro H. Milet, and Nicolau C. Saldanha. On the connectivity of spaces of three-dimensional tilings. *Transactions of the American Mathematical Society*, 375:1579–1605, 2022.
- [Fou96] J.C. Fournier. Pavage des figures planes sans trous par des dominos: Fondement graphique de l’algorithme de Thurston, parallélisation, unicité et décomposition. *Theoretical Computer Science*, 159(1):105–128, 1996. Selected Papers from the “GASCOM ’94” and the “Polynominoes and Tilings” workshops.
- [Geo11] Hans-Otto Georgii. *Gibbs measures and phase transitions*, volume 9 of *De Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, second edition, 2011.
- [GMT20] Alessandro Giuliani, Vieri Mastropietro, and Fabio Lucio Toninelli. Non-integrable dimers: universal fluctuations of tilted height profiles. *Comm. Math. Phys.*, 377(3):1883–1959, 2020.
- [Gor21] Vadim Gorin. *Lectures on random lozenge tilings*, volume 193. Cambridge University Press, 2021.
- [GRT22] Alessandro Giuliani, Bruno Renzi, and Fabio Toninelli. Weakly non-planar dimers. *arXiv preprint arXiv:2207.10428*, 2022.
- [Hal35] P. Hall. On representatives of subsets. *Journal of the London Mathematical Society*, s1-10(1):26–30, 1935.
- [HKMS03] David A. Huse, Werner Krauth, R. Moessner, and S. L. Sondhi. Coulomb and liquid dimer models in three dimensions. *Phys. Rev. Lett.*, 91:167004, Oct 2003.
- [HLT23] Ivailo Hartarsky, Lyuben Lichev, and Fabio Toninelli. Local dimer dynamics in higher dimensions. *arXiv preprint arXiv:2304.10930*, 2023.
- [JPS98] William Jockusch, James Propp, and Peter Shor. Random domino tilings and the Arctic circle theorem. *arXiv preprint arXiv:math/9801068*, 1998.

- [Kel98] Gerhard Keller. *Equilibrium states in ergodic theory*, volume 42. Cambridge University Press, 1998.
- [Ken00] Richard Kenyon. Conformal invariance of domino tiling. *Annals of probability*, pages 759–795, 2000.
- [Ken09] Richard Kenyon. Lectures on dimers. *aXiv preprint arXiv:0910.3129*, 2009.
- [KM03] Werner Krauth and R. Moessner. Pocket Monte Carlo algorithm for classical doped dimer models. *Phys. Rev. B*, 67:064503, Feb 2003.
- [KMT20] Andrew Krieger, Georg Menz, and Martin Tassy. Deducing a variational principle with minimal *a priori* assumptions. *Electron. J. Combin.*, 27(4):Paper No. 4.1, 53, 2020.
- [KOS06] Richard Kenyon, Andrei Okounkov, and Scott Sheffield. Dimers and amoebae. *Annals of mathematics*, pages 1019–1056, 2006.
- [KS22] Caroline J. Klivans and Nicolau C. Saldanha. Domino tilings and flips in dimensions 4 and higher. *Algebr. Comb.*, 5(1):163–185, 2022.
- [Kuc22] Nikolai Kuchumov. A variational principle for domino tilings of multiply connected domains. *arXiv preprint arXiv:2110.06896*, 2022.
- [Lam21] Piet Lammers. A generalisation of the honeycomb dimer model to higher dimensions. *Ann. Probab.*, 49(2):1033–1066, 2021.
- [Lit75] C.H.C Little. A characterization of convertible $(0, 1)$ -matrices. *Journal of Combinatorial Theory, Series B*, 18(3):187–208, 1975.
- [LMN01] Joakim Linde, Christopher Moore, and Mats G. Nordahl. An n -dimensional generalization of the rhombus tiling. In *Discrete models: combinatorics, computation, and geometry (Paris, 2001)*, Discrete Math. Theor. Comput. Sci. Proc., AA, pages 023–042. Maison Inform. Math. Discrèt. (MIMD), Paris, 2001.
- [LR69] O. E. Lanford, III and D. Ruelle. Observables at infinity and states with short range correlations in statistical mechanics. *Comm. Math. Phys.*, 13:194–215, 1969.
- [LZ12] Jeffrey C. Lagarias and Chuanming Zong. Mysteries in packing regular tetrahedra. *Notices Amer. Math. Soc.*, 59(11):1540–1549, 2012.
- [Mil15] Pedro H. Milet. Domino tilings of three-dimensional regions. *PhD thesis, arXiv preprint arXiv:1503.04617*, 2015.
- [MS] Pedro H. Milet and Nicolau C. Saldanha. Enumeration of tilings for the $4 \times 4 \times 4$ box. <http://mat.puc-rio.br/~nicolau/multiplex/example444.html>.
- [MS14a] Pedro H. Milet and Nicolau C. Saldanha. Domino tilings of three-dimensional regions: flips, trits and twists. *arXiv preprint arXiv:1410.7693*, 2014.
- [MS14b] Pedro H. Milet and Nicolau C. Saldanha. Twists for duplex regions. *arXiv preprint arXiv:1411.1793*, 2014.

- [MS15] Pedro H. Milet and Nicolau C. Saldanha. Flip invariance for domino tilings of three-dimensional regions with two floors. *Discrete Comput. Geom.*, 53(4):914–940, 2015.
- [PR14] Benedetto Piccoli and Francesco Rossi. Generalized Wasserstein distance and its application to transport equations with source. *Arch. Ration. Mech. Anal.*, 211(1):335–358, January 2014.
- [PR16] Benedetto Piccoli and Francesco Rossi. On properties of the generalized Wasserstein distance. *Arch. Ration. Mech. Anal.*, 222(3):1339–1365, December 2016.
- [PRT19] Benedetto Piccoli, Francesco Rossi, and Magali Tournus. A Wasserstein norm for signed measures, with application to nonlocal transport equation with source term. *arXiv preprint arXiv:1910.05105*, 2019.
- [QT22] Alexandra Quitmann and Lorenzo Taggi. Macroscopic loops in the 3d double-dimer model. *arXiv preprint arXiv:2206.08284*, 2022.
- [QT23] Alexandra Quitmann and Lorenzo Taggi. Macroscopic loops in the Bose gas, spin $O(n)$ and related models. *Communications in Mathematical Physics*, pages 1–56, 2023.
- [RST99] Neil Robertson, P. D. Seymour, and Robin Thomas. Permanents, Pfaffian orientations, and even directed circuits. *Annals of Mathematics*, 150(3):929–975, 1999.
- [RT00] Dana Randall and Prasad Tetali. Analyzing Glauber dynamics by comparison of Markov chains. *Journal of Mathematical Physics*, 41(3):1598–1615, 2000.
- [RY00] Dana Randall and Gary Yngve. Random three-dimensional tilings of Aztec octahedra and tetrahedra: an extension of domino tilings. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (San Francisco, CA, 2000)*, pages 636–645. ACM, New York, 2000.
- [Sal20] Nicolau C. Saldanha. Domino tilings in dimension 3. In *2020 Fall Central Sectional Meeting*. AMS, 2020.
- [Sal21] Nicolau C. Saldanha. Domino tilings of cylinders: connected components under flips and normal distribution of the twist. *Electron. J. Combin.*, 28(1):Paper No. 1.28, 23, 2021.
- [Sal22] Nicolau C. Saldanha. Domino tilings of cylinders: the domino group and connected components under flips. *Indiana Univ. Math. J.*, 71(3):965–1002, 2022.
- [She05] Scott Sheffield. Random surfaces. *Astérisque*, 304, 2005.
- [Str10] Gilbert Strang. Maximum flows and minimum cuts in the plane. *J. Global Optimization*, 47:527–535, 07 2010.
- [Tag22] Lorenzo Taggi. Uniformly positive correlations in the dimer model and macroscopic interacting self-avoiding walk in \mathbb{Z}^d , $d \geq 3$. *Communications on Pure and Applied Mathematics*, 75(6):1183–1236, 2022.

- [Thu90] William P. Thurston. Conway's tiling groups. *The American Mathematical Monthly*, 97(8):757–773, 1990.
- [Var16] S.R.S. Varadhan. *Large deviations*. Courant Lecture Notes. Courant Institute of Mathematical Sciences, 2016.
- [vdB93] J. van den Berg. A uniqueness condition for Gibbs measures, with application to the 2-dimensional Ising antiferromagnet. *Comm. Math. Phys.*, 152(1):161–166, 1993.
- [vdBS94] J. van den Berg and J. E. Steif. Percolation and the hard-core lattice gas model. *Stochastic Process. Appl.*, 49(2):179–197, 1994.
- [Vil09] Cédric Villani. *The Wasserstein distances*, pages 93–111. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009.
- [WMDB02] M. Widom, R. Mosseri, N. Destainville, and F. Bailly. Arctic octahedron in three-dimensional rhombus tilings and related integer solid partitions. *J. Statist. Phys.*, 109(5-6):945–965, 2002.

Index of symbols

Chapter 2

- E - edges in \mathbb{Z}^3 , 21
- $H(\nu)$ - Shannon entropy of the measure ν , 24
- $H_\Lambda(\mu)$, 24
- $X(\sigma)$, 24
- $\text{Ent}(g)$ - entropy of an asymptotic flow g , 25
- $\Omega(\Lambda)$ - dimer tilings of \mathbb{Z}^3 restricted to a finite set Λ , 24
- Ω - the space of dimer tilings, 20
- \mathcal{P}^s - the space of probability measures with mean current s , 24
- $\text{ent}(s)$ - entropy of the mean current s , 25
- η_i and e_i - unit vectors and edges connecting the origin, 21
- \mathcal{O} - mean current octahedron/possible values of the mean current, 24
- \mathcal{P} - the space of $\mathbb{Z}_{\text{even}}^3$ invariant measures on the space of tilings, 22
- $\mathcal{P}_e, \mathcal{P}_G, \mathcal{P}_{G,e}$ - ergodic, Gibbs and ergodic Gibbs measures respectively on the space of tilings, 23
- $\text{div } v$ - divergence of a discrete vector field v , 21
- $\mathbb{Z}_{\text{even}}^3$, 22
- $f_{(\tau_1, \tau_2)}$ - discrete flow associated to a double dimer configuration, 22
- f_τ - (divergence free) tiling flow associated with the tiling τ , 21
- $h(\mu)$ - specific entropy of the measure μ , 24
- r - reference flow, 21
- $s(\mu)$ - the mean current of the measure μ , 23
- $s_0(\tau)$ - the direction of the tile in τ at the origin, 24
- $s_n(\tau)$ - the average direction of tiles in τ in the box Λ_n , 24
- v_τ - flow associated to a tiling τ , 21

Chapter 4

- ent_{loz} - the entropy function for ergodic Gibbs measures of a given slope on the space of lozenge tilings, 40
- $A_n(c), A_{n,n}$, 38
- L_c - a slab in \mathbb{Z}^3 , 35

$\mathbb{Z}_{\text{loz}}, \mathcal{P}_{\text{loz}}$ - the space of invariant probability measures on the space of lozenge tilings/ dimer tilings of the hexagonal lattice, 36

\mathcal{O}_d the space of admissible mean currents in d dimensions, 33

ρ_0 - marginal of a measure μ to the slab L_0 , 37

τ_B - the restriction of a tiling τ to B , 37

τ_c - the tiling τ on the slab L_c , 37

g_ρ - extreme Gibbs decomposition, 37

$s(\mu)$ - an abuse of notation specific to this section where it also represents the slope of lozenge tilings, 36

Chapter 5

$AF(R)$ - the space of asymptotic flows, 43

$AF^\infty(R)$ - the space of smooth asymptotic flows on R , 51

R - domain, 42

R_ϵ - the domain R with an ϵ boundary removed, 51

$T : AF^\infty(R) \times \mathbb{S}(R) \rightarrow \mathcal{M}^s(R)$ - the trace operator on smooth asymptotic flows, 52

$T : AF(R) \times \mathbb{S}(R) \rightarrow \mathcal{M}^s(R)$ - the trace operator on asymptotic flows, 56

$T : TF(R) \times \mathbb{S}(R) \rightarrow \mathcal{M}^s(R)$ - trace operator for tiling flows, 57

$TF_n(R)$ - free-boundary tiling flows on R at scale n , 42

$T_n(R)$ - free boundary tilings at scale n , 42

$\mathbb{S}(R)$ - compact piecewise smooth surfaces contained in R , 52

$\mathbb{W}_1^{1,1}$ - generalized Wasserstein distance between signed measures, 44

$\mathcal{M}(R), \mathcal{M}^s(R), M_{\text{ac}}^s(R, a, b), \mathcal{M}_{\text{ac}}^s(R)$ - Various spaces of measures on R , 45

∂R - boundary of the domain, 42

τ_R - collection of tiles intersecting R , 42

d_W - Wasserstein distance on flows, 45

Chapter 6

B_n - the box $[-n, n]^3$ and the annulus

$A_n = B_n \setminus B_{(1-\delta)n}$, 68

$\text{Cov}_\epsilon(X)$ - ϵ covering area, 76

flux - flux of a discrete vector field, 69

$\text{flux}(v_\tau, S)$ - flux of the flow v_τ through the surface S , 78

imbalance, 66

set C , its neighbors $N(C)$, 65

Chapter 7

C_P, A_P , 97

D_P , 100

$F : \Omega \times \Omega \rightarrow [-2, 2]^3$, 91

F^* - “slope function”, 91

$S_{\text{first}}, X_{\text{last}}$, 99

\mathcal{P}^{s_1, s_2} , - the space of probability measures on the double dimer model with mean-currents s_1 and s_2 and the two projections π_1, π_2 , 88

\mathcal{E} the edges of the mean-current octahedron \mathcal{O} , 84

$d_P(\alpha)$, 98

given a measure μ with marginals μ_1, μ_2 on double dimer configurations, the swapped measure is usually denoted by μ' with marginals μ'_1, μ'_2 , 93

given a pair of tilings (τ_1, τ_2) , the swapped tiles are usually represented by (τ'_1, τ'_2) , 93

Chapter 8

$(\theta_n)_{n \geq 1}$ - admissible threshold sequence, 114

$A_\delta(g)$ - Wasserstein open balls of radius δ around g , 115

I_b - rate function, 114

\mathcal{X} - tetrahedral mesh, 119

$\bar{\rho}_n, \bar{\mu}_n, \bar{Z}_n$, 118

ρ_n, μ_n, Z_n , 114

τ_v , 122

\tilde{g} - usually denotes the piecewise approximation of g , 119

$v_n = n^3 \text{Vol}(R)$ is roughly the number of vertices of $\frac{1}{n} \mathbb{Z}^3$ in R , 114

$w(v)$, 123

Index of terms

Chapter 1

brickwork tilings/pattern, 4

Chapter 2

discrete vector field/discrete flow, 21
double dimer tiling/configuration, 22
EGM/ergodic Gibbs measures, 23
entropy of an asymptotic flow, 25
ergodic components/decomposition, 23
ergodic measures, 23
Gibbs measures, 23
local/mean-current entropy function, 25
mean current, 23
mean current octahedron, 24
pretiling flow, 21
Shannon entropy, 24
specific entropy, 24
tiling flow, 21

Chapter 3

hopfion, 26
local move/flip, 26
loop, 30
loop shift, 30
trit, 26
twist invariant, 27

Chapter 4

extreme Gibbs measures, 37
frozen tilings, 34
north, east and up tiles, 35
slabs, 35
slope for lozenge tilings, 36

Chapter 5

asymptotic flow, 43
domain, 42
free boundary tiling flows on R at scale n , 42
free boundary tilings of R , 42
free boundary tilings of R at scale n , 42
generalized Wasserstein distance, 44
grid region of scale n , 42
measurable divergence-free vector field, 43
measures corresponding to a free-boundary tiling flow, 46
measures corresponding to a measurable vector field, 45
scale n grid regions contained in domains:
 $R \subset R_n$, 42

trace operator for asymptotic flows, 56
trace operator for tiling flows, 57
trace operator smooth asymptotic flows, 52

Wasserstein distance on flows, 45

Chapter 6

ϵ covering area, 76
 ϵ patch, 69
 ϵ -nearly-constant, 69
a cut, 65
balanced, 64
black and white surface, 66
counterexample to tileability, 66
discrete surface, 66
flux, 69, 78
imbalance, 66
minimal (monochromatic) surface, 67
monochromatic surface, 66

Chapter 7

chain swapping, swapped measure, 93
frozen points in a domain, 108
semi-flexible, rigid, 108

Chapter 8

admissible threshold sequence, 114
boundary conditions which are extendable outside, 113
channel, 126
flexible, 116
hard boundary large deviation principle (HB LDP), 117
periodic tiling, 122
piecewise constant approximation, 119
soft boundary large deviation principle (SB LDP), 114
tetrahedral mesh, 119